COGNITIVE DIGITAL TWIN FRAMEWORK: MODELING AND REAL-TIME DECISION MAKING

Yangyang Zhang[1], **Mengtong Li**[2], **Xinyu Wang**[3,†], **Zhihao Lin**[4], **Xiang Luo**[3], **Zhen Tian**[4], **Ning Lyu**[5], **Zhiguo Tao**[6], **Aaron Wang**[7]

[1]JAC Motors Co., Ltd., China [2]Siemens Energy Innovation Center [3]Hefei University of Technology, China
[4]University of Glasgow, UK [5]Carnegie Mellon University, USA
[6]University of Nottingham, UK [7]Tongji University, China

ABSTRACT

Digital twins can support modeling and optimization of complex physical systems. However, many existing frameworks remain passive and cannot learn or decide online under non-stationary conditions. We present a Cognitive Digital Twin (CDT) framework that couples real-time perception, reasoning, and adaptation. CDT builds on deep reinforcement learning, temporal knowledge graphs, and federated learning. It makes three design choices. First, it uses a quality-aware multi-modal fusion module to weight heterogeneous inputs. Second, it uses a hierarchical reasoning engine with reactive (0–10 ms), deliberative (10 ms–1 s), and reflective (asynchronous) layers. Third, it uses a privacy-preserving federated protocol to coordinate multiple twins. Experiments on three industrial IoT testbeds show that CDT reduces prediction error by 25.5%, reduces average response latency by 30.0%, and improves decision quality by 25.7% over the best baseline. Ablations show that each module contributes, and the full system gives the best overall performance.

1 INTRODUCTION

Digital twins enable synchronized cyber–physical representations for monitoring and optimization of physical systems (Barricelli et al., 2019; Fuller et al., 2020; Tao et al., 2019). However, most existing digital twins remain largely passive, focusing on visualization, simulation, and offline analytics, and often fail to support reliable online decision-making in dynamic, data-intensive environments (Minerva et al., 2020; Semeraro et al., 2021; Hribernik et al., 2021; Liu et al., 2024).

The proliferation of IoT devices and edge platforms has introduced large-scale, heterogeneous, and non-stationary data streams (Tao et al., 2019; Minerva et al., 2020), placing stringent requirements on real-time perception, adaptation, and control (Jeremiah et al., 2024). In practice, current digital-twin architectures commonly exhibit limited online adaptability, heavy reliance on preset or human-in-the-loop decisions, weak coordination across distributed twins, and insufficient semantic reasoning for causal consistency and knowledge reuse (Jones et al., 2020; Jeremiah et al., 2024).

To address these challenges, we propose a *Cognitive Digital Twin (CDT)* framework for online perception and decision-making (Lu et al., 2022; Eirinakis et al., 2022). CDT integrates deep reinforcement learning for adaptive control (Kim et al., 2020; Lee et al., 2022), temporal knowledge graphs for semantic context and constraint validation (Stavropoulou et al., 2024), and privacy-preserving federated learning for cross-twin collaboration without raw-data sharing (Kairouz et al., 2021; Li et al., 2020). Beyond combining these components, CDT introduces a time-budgeted memory mechanism with *write* and *retrieve* operations, linking episodic experience and symbolic knowledge to layered decision processes.

Although attention-based fusion, knowledge graphs, and federated learning have been individually explored, their joint deployment in real-time cognitive twins introduces coupled constraints on latency, consistency, and coordination. CDT addresses these issues through a memory-centric,

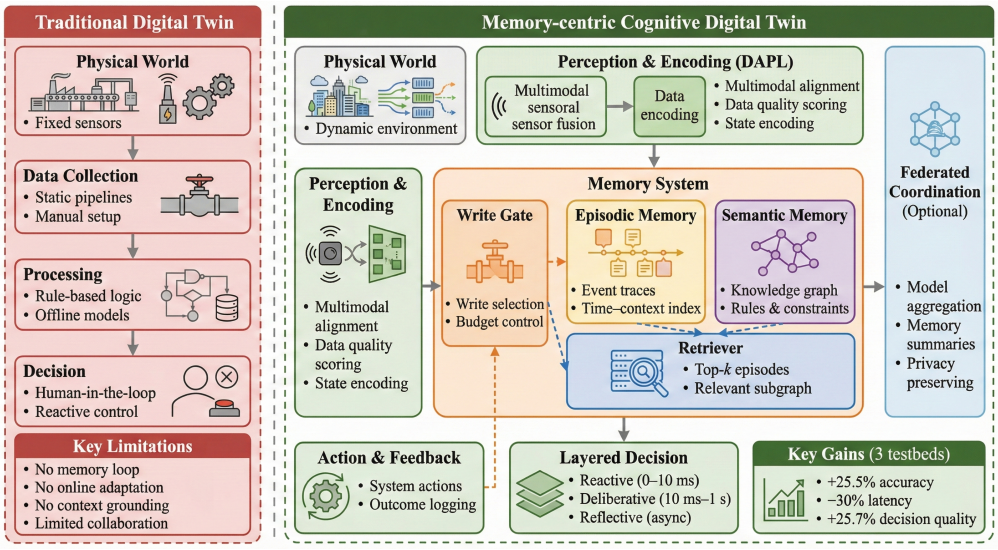

Figure 1: Conceptual comparison between a traditional digital twin and the proposed CDT. CDT augments perception with time-budgeted memory (write/retrieve) and supports layered decisions under real-time constraints.

time-budgeted architecture that explicitly aligns perception reliability, semantic grounding, and hierarchical reasoning.

The main contributions are summarized as follows:

- **Quality-aware perception.** A reliability-aware multi-modal fusion module that produces stable state representations from heterogeneous inputs.
- **Memory-centric hierarchical reasoning.** A write/retrieve memory system coupled with a three-level reasoning engine spanning reactive, deliberative, and reflective time scales.
- **Federated coordination and evaluation.** A privacy-preserving federated protocol evaluated on three industrial IoT testbeds with comprehensive results and ablations.

The remainder of this paper is organized as follows. Section 2 presents the CDT formulation and design. Section 3 reports experimental results. Section 4 concludes the paper, with related work discussed in Appendix A.

2 METHODOLOGY

2.1 PROBLEM FORMULATION

We model CDT as a partially observable Markov decision process (POMDP):

$$\mathcal{M} = \big(\mathcal{S}, \mathcal{O}, \mathcal{A}, \mathcal{T}, R, \gamma\big), \tag{1}$$

where $\mathcal{S}$ is the latent state space, $\mathcal{O}$ is the observation space, and $\mathcal{A}$ is the action space. The transition kernel is $\mathcal{T} : \mathcal{S} \times \mathcal{A} \to \Delta(\mathcal{S})$, where $\Delta(\mathcal{S})$ denotes the set of probability distributions over $\mathcal{S}$. The reward function is $R : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$, and the discount factor is $\gamma \in (0, 1)$.

At time step t, the system occupies a latent state $\mathbf{s}_t \in \mathcal{S}$. The agent receives an observation $\mathbf{o}_t \in \mathcal{O}$ through an emission function $h : \mathcal{S} \to \mathcal{O}$ with additive Gaussian noise:

$$\mathbf{o}_t = h(\mathbf{s}_t) + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}), \tag{2}$$

where $\boldsymbol{\Sigma} \in \mathbb{R}^{|\mathcal{O}| \times |\mathcal{O}|}$ is the observation noise covariance matrix.

CDT operates under time-varying constraints that capture safety, resource budgets, and real-time latency. The agent selects an action $\mathbf{a}_t \in \mathcal{A}_t$ from the feasible set

$$\mathcal{A}_t = \big\{\mathbf{a} \in \mathcal{A} \;\big|\; g_j(\mathbf{s}_t, \mathbf{a}) \le 0,\ \mathbf{c}^\top \mathbf{a} \le C_{\max},\ \tau(\mathbf{a}) \le \tau_{\max}\big\}, \tag{3}$$

where $g_j : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ is the j-th safety constraint, $\mathbf{c} \in \mathbb{R}^{|\mathcal{A}|}$ is a cost-weight vector with budget $C_{\max}$, and $\tau : \mathcal{A} \to \mathbb{R}_{\ge 0}$ measures execution latency with limit $\tau_{\max}$.

To reflect the memory-centric design, we introduce an internal memory state $\mathbf{m}_t \in \mathcal{M}_{\text{mem}}$ that summarizes the observation history under a time budget. The memory is updated by a write operator

$$\mathbf{m}_{t+1} = \text{Write}(\mathbf{m}_t, \mathbf{o}_t), \tag{4}$$

and the policy conditions on the current observation and memory:

$$\pi(\mathbf{a}_t \mid \mathbf{o}_t, \mathbf{m}_t). \tag{5}$$

We consider K objectives and use a weighted scalarization:

$$R(\mathbf{s}_t, \mathbf{a}_t) = \sum_{k=1}^{K} w_k\, R_k(\mathbf{s}_t, \mathbf{a}_t), \quad w_k \ge 0,\ \sum_{k=1}^{K} w_k = 1, \tag{6}$$

where each $R_k : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$ encodes a distinct goal (e.g., efficiency, safety, and quality).

The goal is to find an optimal policy maximizing the expected discounted return over horizon H:

$$\pi^* = \underset{\pi}{\text{argmax}}\ \mathbb{E}_\pi\left[\sum_{t=0}^{H} \gamma^t\, R(\mathbf{s}_t, \mathbf{a}_t)\right], \quad \text{s.t. } \mathbf{a}_t \in \mathcal{A}_t,\ \forall t. \tag{7}$$

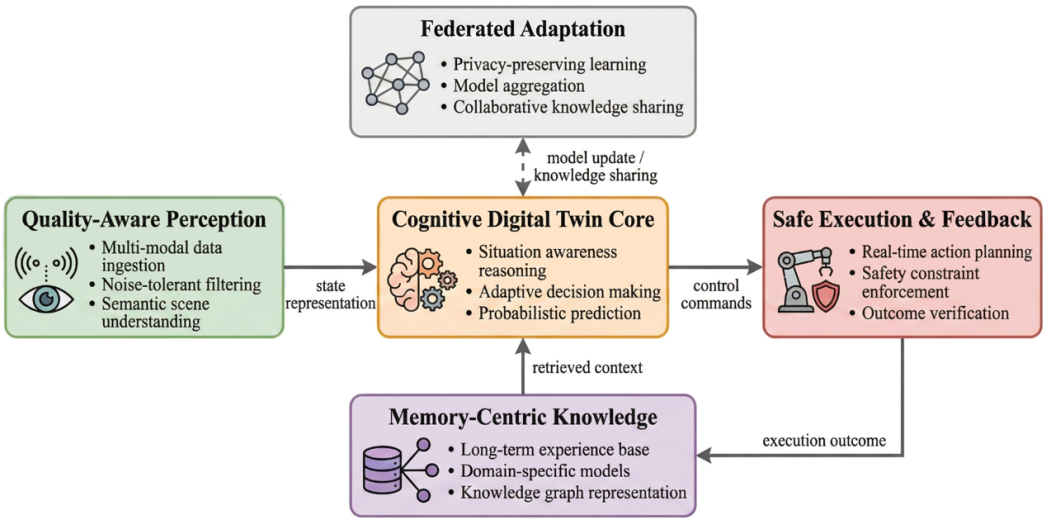

Figure 2: CDT architecture aligned with Sec. 2: DAPL performs quality-aware ingestion and alignment of heterogeneous data; CPE fuses multi-modal features via attention; KMS maintains a temporal knowledge graph for semantic reasoning and constraint validation; HDMM is shown as a conceptual downstream decision interface.

2.2 Cognitive Digital Twin Architecture

The CDT framework comprises five modules organized in a layered architecture (Fig. 2): Data Acquisition and Preprocessing Layer (DAPL), Cognitive Processing Engine (CPE), Knowledge Management System (KMS), Hierarchical Decision-Making Module (HDMM), and Federated Learning Coordination Layer (FLCL).

2.2.1 Data Acquisition and Preprocessing Layer (DAPL)

DAPL processes heterogeneous inputs—sensor streams, images, time series, and telemetry—through modality-specific pipelines for alignment, denoising, and validation (Minerva et al., 2020; Liu et al., 2024). Each sample $\mathbf{x}$ is assigned a composite reliability score $q(\mathbf{x}) \in [0, 1]$:

$$q(\mathbf{x}) = \alpha\, q_{\text{comp}}(\mathbf{x}) + \beta\, q_{\text{cons}}(\mathbf{x}) + \gamma\, q_{\text{time}}(\mathbf{x}), \tag{8}$$

where $\alpha, \beta, \gamma \in [0, 1]$ and $\alpha + \beta + \gamma = 1$.

The completeness term penalizes missing or invalid fields:

$$q_{\text{comp}}(\mathbf{x}) = 1 - \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}_{\text{invalid}}(x_i), \tag{9}$$

where x_i denotes the i-th scalar field in $\mathbf{x}$, N is the number of fields, and $\mathbb{I}_{\text{invalid}}(\cdot) \in \{0, 1\}$ is an indicator of invalidity.

The consistency term aggregates range, logic, and historical checks:

$$q_{\text{cons}}(\mathbf{x}) = \phi_{\text{range}}(\mathbf{x})\, \phi_{\text{logic}}(\mathbf{x})\, \phi_{\text{hist}}(\mathbf{x}), \tag{10}$$

where $\phi_{\text{range}}, \phi_{\text{logic}}, \phi_{\text{hist}} \in [0, 1]$ are scalar validation functions.

The timeliness term decays exponentially with data latency Δt:

$$q_{\text{time}}(\mathbf{x}) = \exp(-\lambda\, \Delta t)\,, \tag{11}$$

where $\Delta t \geq 0$ (unit: s) and $\lambda > 0$ (unit: s^{-1}) controls the decay rate. The resulting score $q(\mathbf{x})$ enables quality-aware downstream fusion.

2.2.2 COGNITIVE PROCESSING ENGINE (CPE)

CPE employs modality-specific encoders $f_m : \mathcal{X}_m \to \mathbb{R}^d$ that map heterogeneous inputs into a shared d-dimensional latent space:

$$\mathbf{f}_m = f_m(\mathbf{X}_m), \quad m \in \{\text{sensor, image, seq, text}\}, \tag{12}$$

where $\mathbf{f}_m \in \mathbb{R}^d$ is the modality feature vector and $\mathbf{X}_m \in \mathcal{X}_m$ is the corresponding input.

Multi-modal fusion is performed via quality-aware attention. The attention weight for modality m is computed as

$$\alpha_m = \frac{\exp(s_m)}{\sum_j \exp(s_j)}, \quad s_m = \frac{\mathbf{u}_m^\top \mathbf{v}_m}{\sqrt{d_k}} + \log\big(\max(q_m, \epsilon)\big), \tag{13}$$

where $\mathbf{u}_m, \mathbf{v}_m \in \mathbb{R}^{d_k}$ are query and key vectors, d_k is the attention dimension, $q_m \in [0, 1]$ is the reliability score from DAPL for modality m, and $\epsilon > 0$ is a numerical stabilizer. The fused representation is

$$\mathbf{f} = \sum_m \alpha_m \, \mathbf{f}_m. \tag{14}$$

2.2.3 KNOWLEDGE MANAGEMENT SYSTEM (KMS)

KMS maintains a temporal knowledge graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{R})$ with entity set $\mathcal{V}$, edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{R} \times \mathcal{V}$, and relation types $\mathcal{R}$. Using temporal TransE-style embeddings, the plausibility of a timestamped quadruple (u, r, v, τ) is scored as

$$s(u, r, v, \tau) = -\left\| \mathbf{e}_u + \mathbf{e}_r + \boldsymbol{\psi}(\tau) - \mathbf{e}_v \right\|_2, \tag{15}$$

where $\mathbf{e}_u, \mathbf{e}_v \in \mathbb{R}^d$ are head and tail entity embeddings, $\mathbf{e}_r \in \mathbb{R}^d$ is the relation embedding, and $\boldsymbol{\psi} : \mathbb{R} \to \mathbb{R}^d$ encodes the timestamp τ. Graph updates are triggered when confidence exceeds a threshold, while stale knowledge is attenuated via temporal decay.

2.3 HIERARCHICAL COGNITIVE REASONING ENGINE

CDT employs a three-level hierarchical reasoning engine with explicit time budgets, balancing real-time responsiveness with long-horizon adaptation.

Reactive layer (0–10 ms). This layer ensures immediate safety through a lightweight emergency detector:

$$\boldsymbol{a}_t^{\text{react}} = \begin{cases} \boldsymbol{a}^{\text{stop}}, & p_{\text{crit}} > 0.9, \\ \boldsymbol{a}^{\text{shed}}, & p_{\text{warn}} > 0.7, \\ \boldsymbol{a}^{\text{adj}}, & p_{\text{anom}} > 0.5, \\ \boldsymbol{a}^{\text{keep}}, & \text{otherwise}, \end{cases} \tag{16}$$

where $\boldsymbol{a}^{\text{stop}}, \boldsymbol{a}^{\text{shed}}, \boldsymbol{a}^{\text{adj}}, \boldsymbol{a}^{\text{keep}} \in \mathcal{A}$ denote the predefined action primitives for emergency stop, load shedding, parameter adjustment, and normal operation, respectively.

Deliberative layer (10 ms–1 s). When sufficient time is available, this layer performs short-horizon planning using a hybrid predictor that combines physics-based models with learned residual corrections:

$$\hat{\mathbf{s}}_{t+1} = f_{\text{phy}}(\mathbf{s}_t, \mathbf{a}_t) + f_\Delta(\mathbf{s}_t, \mathbf{a}_t; \boldsymbol{\theta}), \tag{17}$$

where $f_{\text{phy}} : \mathcal{S} \times \mathcal{A} \to \mathcal{S}$ encodes known physical dynamics, f_Δ is a neural residual model parameterized by $\boldsymbol{\theta}$, and $\hat{\mathbf{s}}_{t+1} \in \mathcal{S}$ is the one-step predicted state. The layer then solves a model predictive control (MPC) problem over planning horizon H.

Reflective layer (asynchronous). Operating outside the critical path, this layer leverages prioritized experience replay and model-agnostic meta-learning (MAML) to adapt long-term strategies based on accumulated experience. Algorithm 1 summarizes the integrated reasoning process.

Here t_{avail} denotes the estimated time budget for deliberation, $\mathbf{g}_t$ denotes the goal set queried from $\mathcal{G}$, and $\mathcal{D}_{\text{replay}}$ is the prioritized replay buffer used by the reflective layer.

Algorithm 1 Hierarchical Cognitive Reasoning (Compact)

Input: observation $\boldsymbol{o}_t$, context $\boldsymbol{c}_t$, graph $\mathcal{G}$, state $\boldsymbol{s}_t$
Output: action $\boldsymbol{a}_t$, updated graph $\mathcal{G}'$
$(p_{\text{crit}}, p_{\text{warn}}, p_{\text{anom}}) \leftarrow f_{\text{risk}}(\boldsymbol{o}_t)$
if $p_{\text{crit}} > 0.9$ **then**
  return $\boldsymbol{a}_t \leftarrow \boldsymbol{a}^{\text{stop}}$
else if $p_{\text{warn}} > 0.7$ **then**
  return $\boldsymbol{a}_t \leftarrow \boldsymbol{a}^{\text{shed}}$
else if $p_{\text{anom}} > 0.5$ **then**
  return $\boldsymbol{a}_t \leftarrow \boldsymbol{a}^{\text{adj}}$
end if
if $t_{\text{avail}} > 10\,\text{ms}$ **then**
  $\boldsymbol{g}_t \leftarrow \text{QueryKG}(\mathcal{G}, \boldsymbol{s}_t, \boldsymbol{c}_t)$
  $\hat{\boldsymbol{s}}_{t+1:t+H} \leftarrow \text{Predict}(\boldsymbol{s}_t, \boldsymbol{g}_t, H)$
  $\boldsymbol{a}_t \leftarrow \text{SolveMPC}(\boldsymbol{s}_t, \boldsymbol{g}_t, \hat{\boldsymbol{s}}_{t+1:t+H})[0]$
else
  $\boldsymbol{a}_t \leftarrow \pi_{\text{react}}(\boldsymbol{o}_t, \boldsymbol{c}_t)$
end if
if $t \bmod T_{\text{reflect}} = 0$ **then**
  $\boldsymbol{\theta} \leftarrow \text{MAMLUpdate}(\boldsymbol{\theta}, \mathcal{D}_{\text{replay}})$
  $\mathcal{G}' \leftarrow \text{UpdateKG}(\mathcal{G}, \boldsymbol{\theta})$
else
  $\mathcal{G}' \leftarrow \mathcal{G}$
end if
return $\boldsymbol{a}_t, \mathcal{G}'$

3 EXPERIMENTS

3.1 EXPERIMENTAL SETUP

We evaluate CDT on three industrial IoT testbeds spanning manufacturing, buildings, and logistics. Unless otherwise specified, all methods share the same preprocessing and data split (70/15/15). Hyperparameters are selected on the validation set with early stopping (patience $=$ 10 epochs). Results are reported as mean$\pm$std over five random seeds.

Smart Manufacturing. We adopt an OpenAI Gym-based Industry 4.0 simulation with 15 networked assets (five CNC machines, four assembly robots, three inspection stations, two conveyors, and facility-level sensors), with $>$200 sensing channels (Bi et al., 2024). Data streams include multi-rate time series sampled at 1–1000 Hz and 1080p images at 30 FPS, yielding $\sim$2.5 TB/day. We consider predictive maintenance, visual quality inspection, and short-term load forecasting. For load forecasting, we follow the BuildingsBench protocol and use its recommended pretraining/evaluation split when applicable (Emami et al., 2023). Sensor noise and failure modes are injected following profiles from real-world deployments.

Building Automation. We use the public BuildingsBench dataset representing a 10-story commercial building with over 800 IoT devices monitoring HVAC, lighting, occupancy, and energy subsystems. The dataset contains 7,200 hourly samples and is evaluated under the BuildingsBench protocol for short-term load forecasting and anomaly detection (Emami et al., 2023).

Supply Chain Management. We evaluate a synthetic distributed logistics network generated using SimPy discrete-event simulation, with 25 warehouses, 100+ vehicles, and 500+ suppliers processing more than 50,000 daily transactions. The dataset covers eight months (10 M transactions and 2 M logistics events), supporting demand forecasting, route optimization, and resilience analysis.

Baselines. We compare CDT with five digital-twin baselines under matched budgets and identical input modalities: (1) **Traditional DT**: a rule-based twin with static thresholds and expert-crafted decision rules; (2) **ML-DT**: classical machine learning models (Random Forest, SVM, Logistic Regression) with cross-validated hyperparameters; (3) **DL-DT**: deep learning twins using LSTM/CNN/MLP backbones selected by validation performance; (4) **RL-DT**: a DQN-based con-

Table 1: Performance comparison across testbeds. Best results are in bold. Relative improvements are computed *per metric* against the strongest baseline in the corresponding column. ↓: lower is better; ↑: higher is better.

Method	Prediction		Latency (ms)		Decision		Resource
	MAE↓	F1↑	Mean↓	P95↓	Reward↑	Consist.↑	Effic.↑
Traditional DT	0.184	0.623	450	1200	6.2	0.67	0.72
ML-DT	0.167	0.689	420	1050	6.8	0.71	0.75
DL-DT	0.145	0.734	380	890	7.3	0.74	0.78
RL-DT	0.158	0.701	350	820	7.6	0.76	0.76
FL-DT	0.151	0.718	365	865	7.1	0.73	0.79
CDT (Ours)	**0.108**	**0.847**	**315**	**675**	**9.2**	**0.89**	**0.91**
Rel. Improvement	25.5%	15.4%	10.0%	17.7%	21.1%	17.1%	15.2%

trol twin (3-layer MLP, ϵ-greedy exploration, replay buffer 100k) trained with the same interaction budget as CDT; (5) **FL-DT**: a FedAvg-based federated twin (10 clients, 10 local epochs, learning rate 0.01) using the same client partitioning and communication rounds as CDT. All baselines share the same preprocessing and are evaluated on the same held-out test set.

Evaluation metrics. We quantify performance along four dimensions. *(i) Predictive accuracy:* for regression tasks, we report the mean absolute error (MAE)

$$\mathrm{MAE} = \frac{1}{N}\sum_{i=1}^{N} |y_i - \hat{y}_i|\,, \tag{18}$$

where $y_i \in \mathbb{R}$ and $\hat{y}_i \in \mathbb{R}$ denote the ground-truth and prediction, respectively. For classification/anomaly detection, we report macro-F1

$$\mathrm{F1}_{\mathrm{macro}} = \frac{1}{C}\sum_{c=1}^{C} \frac{2\ \mathrm{Prec}_c\ \mathrm{Rec}_c}{\mathrm{Prec}_c + \mathrm{Rec}_c}, \tag{19}$$

where C is the number of classes and $\mathrm{Prec}_c, \mathrm{Rec}_c \in [0,1]$ are the precision and recall for class c.

(ii) Responsiveness: we measure end-to-end decision latency from observation arrival to action output, reporting the sample mean latency $\bar{\tau}$ and the 95th-percentile latency τ_{95} (unit: ms).

(iii) Decision quality: we report cumulative reward $\sum_t r_t$ (normalized to $[0,10]$ per episode) for control tasks, where $r_t \in \mathbb{R}$ is the scalar reward at time step t, combining task completion rate, constraint satisfaction, and resource consumption. We additionally report decision consistency as the agreement rate over repeated trials under matched initial conditions:

$$\mathrm{Cons} = \frac{1}{M}\sum_{m=1}^{M} \mathbb{I}\Big(a^{(m)} = \tilde{a}^{(m)}\Big)\,, \tag{20}$$

where $a^{(m)}$ and $\tilde{a}^{(m)}$ are the selected high-level actions from two runs with the same scenario seed.

(iv) Resource efficiency: we report a normalized efficiency score that penalizes computation while rewarding task performance,

$$\mathrm{Eff} = \frac{\mathrm{Perf}}{\mathrm{Cost}}, \tag{21}$$

where Cost denotes the average compute cost per decision (e.g., wall-clock time or FLOPs) and Perf is the task-specific performance measure (higher is better).

3.2 Main Results

Table 1 and Fig. 3 summarize the aggregated performance across all testbeds. CDT achieves the strongest overall results, with consistent gains in prediction, latency, decision quality, and efficiency.

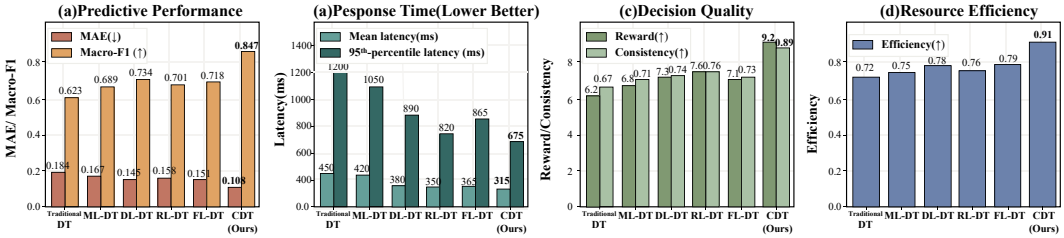

Figure 3: Aggregated performance across four metric categories. CDT attains MAE $= 0.108$ and macro-F1 $= 0.847$, with mean latency $= 315\,\text{ms}$ and $\tau_{95} = 675\,\text{ms}$. It also achieves reward $= 9.2$, consistency $= 0.89$, and resource efficiency $= 0.91$.

Table 2: Ablation results on aggregated testbeds. Values report absolute metrics; percentages indicate relative degradation with respect to the full CDT configuration.

Configuration	MAE↓	F1↑	Latency (ms)↓	Reward↑
Full CDT	0.108	0.847	315	9.2
w/o Multi-modal Fusion	0.127 (+17.6%)	0.789 (-6.8%)	340 (+7.9%)	8.4 (-8.7%)
w/o Hierarchical Reasoning	0.134 (+24.1%)	0.776 (-8.4%)	495 (+57.1%)	7.9 (-14.1%)
w/o Knowledge Graph	0.119 (+10.2%)	0.812 (-4.1%)	350 (+11.1%)	8.7 (-5.4%)
w/o Federated Learning	0.115 (+6.5%)	0.831 (-1.9%)	325 (+3.2%)	8.9 (-3.3%)
w/o Context Awareness	0.122 (+13.0%)	0.798 (-5.8%)	365 (+15.9%)	8.3 (-9.8%)
w/o Adaptive Control	0.111 (+2.8%)	0.839 (-0.9%)	380 (+20.6%)	8.1 (-12.0%)

Predictive accuracy. CDT achieves the lowest MAE (0.108) and the highest macro-F1 (0.847), corresponding to a 25.5% MAE reduction and a 15.4% F1 improvement over the strongest predictive baseline (DL-DT; MAE $= 0.145$, F1 $= 0.734$). The gain is primarily driven by the quality-aware multi-modal fusion, which mitigates unreliable inputs and improves representation robustness. Moreover, integrating the temporal knowledge graph further benefits tasks that require rule-consistent inference and causal attribution (e.g., fault diagnosis).

System responsiveness. CDT reduces the mean end-to-end latency by 10.0% (315 ms vs. 350 ms) and the 95th-percentile latency by 17.7% (675 ms vs. 820 ms) compared with the strongest latency baseline (RL-DT). This improvement aligns with the explicit time budgeting in the hierarchical reasoning engine: the reactive layer supports safety-critical responses within 0–10 ms, while the deliberative layer allocates the remaining budget to short-horizon optimization when feasible. In breakdown, the reactive layer achieves a median latency of 4.2 ms (P95: 8.7 ms), and the deliberative layer averages 285 ms when invoked.

Decision quality. CDT improves cumulative reward by 21.1% (9.2 vs. 7.6) and decision consistency by 17.1% (0.89 vs. 0.76) over the strongest decision baseline (RL-DT). The knowledge graph promotes rule-consistent decisions and reduces inconsistent action switching, while adaptive reinforcement learning helps maintain performance under distribution shift.

3.3 Ablation Studies

Table 2 presents systematic ablations, where we remove one component at a time while keeping the remaining pipeline, training budget, and data splits unchanged. Fig. 4 visualizes the corresponding relative degradation. Unless otherwise stated, *latency* refers to the mean end-to-end inference latency from observation to action.

Multi-modal fusion. Removing multi-modal fusion increases MAE by 17.6% and reduces macro-F1 by 6.8%, indicating that reliability-aware fusion is critical for stabilizing representations under heterogeneous sensing quality.

Hierarchical reasoning. Disabling hierarchical reasoning causes the largest latency increase (+57.1%) and a pronounced reward drop (-14.1%), highlighting the importance of time-budgeted decomposition for balancing rapid response and long-horizon decision quality.

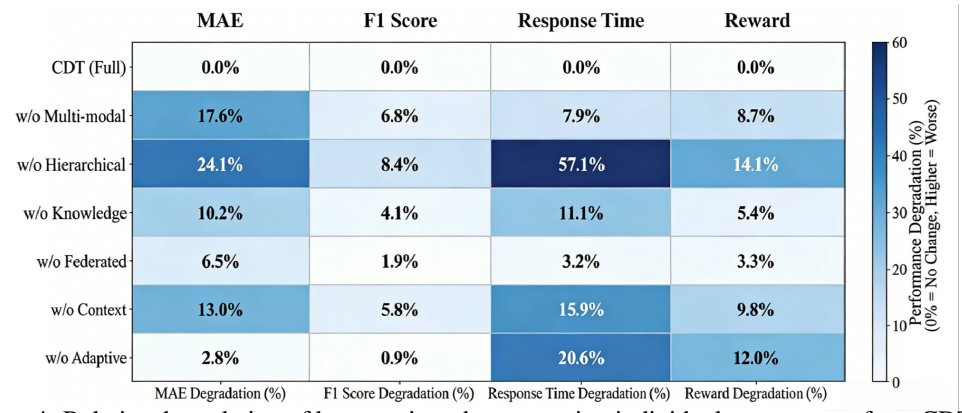

Figure 4: Relative degradation of key metrics when removing individual components from CDT (all changes are computed with respect to the full CDT). Hierarchical reasoning and multi-modal fusion yield the most pronounced degradations.

Knowledge graph reasoning. Removing the knowledge graph degrades MAE by 10.2% and reduces reward by 5.4%, suggesting that explicit semantic constraints and relational context improve prediction consistency and decision coherence.

Federated learning. Federated learning yields modest degradation in centralized evaluation (e.g., +6.5% MAE), but its removal slows cross-site convergence by 25–30% in distributed settings (not shown), underscoring its role in scalable adaptation under non-IID data.

Overall, the ablations show that CDT gains arise from the coupled design of reliability-aware perception, semantic grounding, time-budgeted reasoning, and collaborative learning, rather than from any single component.

4 Conclusion

This paper presented a Cognitive Digital Twin (CDT) framework that advances conventional digital twins from passive monitoring toward closed-loop cognitive agents capable of online perception, reasoning, and decision making. CDT integrates quality-aware multi-modal fusion, temporal knowledge graphs, and time-budgeted hierarchical reasoning to jointly improve prediction accuracy, responsiveness, and decision quality across industrial IoT scenarios. Empirical results and ablation studies demonstrate that these gains arise from the coordinated design of perception, memory, and control, rather than from any individual component.

References

Mohammad A. Al Faruque, Deevakar Muthirayan, Shanglin Y. Yu, and Pramod P. Khargonekar. Cognitive digital twin for manufacturing systems. In *Proceedings of the Design, Automation & Test in Europe Conference (DATE)*, pp. 440–445. IEEE, 2021.

Muhammad Intizar Ali, Pankesh Patel, John G. Breslin, Rafika Harik, and Amit Sheth. Cognitive digital twins for smart manufacturing. *IEEE Intelligent Systems*, 36(2):96–100, 2021. doi: 10.1109/MIS.2021.3062437.

Barbara Rita Barricelli, Elena Casiraghi, and Daniela Fogli. A survey on digital twin: Definitions, characteristics, applications, and challenges. *IEEE Access*, 7:167653–167671, 2019. doi: 10.1109/ACCESS.2019.2953499.

Ziqian Bi, Jiawei Xu, and Ming Liu. Fmcw radar principles and human activity recognition systems: Foundations, techniques, and applications. *arXiv:2410.08483*, 2024.

Zongsheng Cao, Yangfan He, Anran Liu, Jun Xie, Zhepeng Wang, and Feng Chen. Purifygen: A risk-discrimination and semantic-purification model for safe text-to-image generation. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 816–825, 2025.

Elias Dritsas and Maria Trigka. Federated learning for IoT: A survey of techniques, challenges, and applications. *Journal of Sensor and Actuator Networks*, 14(1):9, 2025. doi: 10.3390/jsan14010009.

Pavlos Eirinakis, Stavros Lounis, Stathis Plitsos, George Arampatzis, Kostas Kalaboukas, Klemen Kenda, Jinzhi Lu, Jože M. Rožanec, and Nenad Stojanovic. Cognitive digital twins for resilience in production: A conceptual framework. *Information*, 13(1):33, 2022. doi: 10.3390/info13010033.

Karim El Mokhtari, Ivan Panushev, and J. J. McArthur. Development of a cognitive digital twin for building management and operations. *Frontiers in Built Environment*, 8, 2022. doi: 10.3389/fbuil.2022.856873.

Patrick Emami, Abhijeet Sahu, and Peter Graf. Buildingsbench: A large-scale dataset of 900k buildings and benchmark for short-term load forecasting. *arXiv preprint arXiv:2307.00142*, 2023. URL https://arxiv.org/abs/2307.00142.

Aidan Fuller, Zheng Fan, Charles Day, and Carsten Barlow. Digital twin: Enabling technologies, challenges and open research. *IEEE Access*, 8:108952–108971, 2020. doi: 10.1109/ACCESS.2020.2998358.

Xiangyu Guo, Tang Ji, Yangyang Liu, and Xun Xu. Cognitive digital twin framework for smart manufacturing. In *Proceedings of the 2023 IEEE 19th International Conference on Automation Science and Engineering (CASE)*, pp. 1028–1035, 2023. doi: 10.1109/CASE56687.2023.10260436.

Karl Hribernik, Giacomo Cabri, Federica Mandreoli, and Gregoris Mentzas. Autonomous, context-aware, adaptive digital twins—state of the art and roadmap. *Computers in Industry*, 133:103508, 2021. doi: 10.1016/j.compind.2021.103508.

Sekione R. Jeremiah, Abdellah El Azzaoui, Nengnian Xiong, and Jong Hyuk Park. A comprehensive survey of digital twins: Applications, technologies and security challenges. *Journal of Systems Architecture*, 151:103120, 2024. doi: 10.1016/j.sysarc.2023.103120.

David Jones, Chris Snider, Aydin Nassehi, Jason Yon, and Ben Hicks. Characterising the digital twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology*, 29:36–52, 2020. doi: 10.1016/j.cirpj.2020.02.002.

Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, David Naylor, et al. Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2):1–210, 2021. doi: 10.1561/2200000083.

Young Gun Kim, Sangjun Lee, Ji Son, Hyuk Bae, and Bumseok D. Chung. Multi-agent system and reinforcement learning approach for distributed intelligence in a flexible smart manufacturing system. *Journal of Manufacturing Systems*, 57:440–450, 2020. doi: 10.1016/j.jmsy.2020.10.004.

Timothy Kuo and Hui Yang. Federated learning on distributed and encrypted data for smart manufacturing. *Journal of Computing and Information Science in Engineering*, 24(7):071007, 2024. doi: 10.1115/1.4068980.

Donghoon Lee, Hyojoon Kim, and Hyun-Soo Kim. Digital twin-driven deep reinforcement learning for adaptive task allocation in robotic construction. *Advanced Engineering Informatics*, 53:101710, 2022. doi: 10.1016/j.aei.2022.101710.

Ming Li, Keyu Chen, Ziqian Bi, Ming Liu, Benji Peng, Qian Niu, Junyu Liu, Jinlang Wang, Sen Zhang, Xuanhe Pan, et al. Surveying the mllm landscape: A meta-review of current surveys. *arXiv:2409.18991*, 2024.

Tian Li, Aditya K. Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020. doi: 10.1109/MSP.2020.2975749.

Ruilin Liu, Jihong Yu, Xiaoqing Xie, Stephen Y. Chang, Hongwei Huo, Hing Chau Chan, and Hui Li. A review of digital twin capabilities, technologies, and applications based on the maturity model. *Advanced Engineering Informatics*, 62:102592, 2024. doi: 10.1016/j.aei.2022.102592.

Jinzhi Lu, Zhaorui Yang, Xiaochen Zheng, Jian Wang, and Dimitris Kiritsis. Exploring the concept of cognitive digital twin from model-based systems engineering perspective. *International Journal of Advanced Manufacturing Technology*, 121(9):5835–5854, 2022. doi: 10.1007/s00170-022-09610-5.

Jianhao Lv, Xinyu Li, Yicheng Sun, Yu Zheng, and Jinsong Bao. A bio-inspired lida cognitive-based digital twin architecture for unmanned maintenance of machine tools. *Robotics and Computer-Integrated Manufacturing*, 80:102489, 2023. doi: 10.1016/j.rcim.2022.102489.

Sukanya Mandal. A privacy preserving federated learning (ppfl) based cognitive digital twin (cdt) framework for smart cities. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 23399–23400, 2024.

Sukanya Mandal and Noel E. O'Connor. LLMASMMKG: Llm assisted synthetic multimodal knowledge graph creation for smart city cognitive digital twins. In *Proceedings of the AAAI Symposium Series*, volume 4, pp. 210–221, 2024a.

Sukanya Mandal and Noel E. O'Connor. Ontosmartdcu: A multi-agent microservice based automated, dynamic ontology and knowledge graph creation framework from real-time sensor data for the smart dcu digital twin. In *Proceedings of the 2024 IEEE 10th World Forum on Internet of Things (WF-IoT)*, pp. 784–789. IEEE, 2024b. doi: 10.1109/WF-IoT57584.2024.10062743.

Roberto Minerva, Giacomo M. Lee, and Noel Crespi. Digital twin in the IoT context: A survey on technical features, scenarios, and architectural models. *Proceedings of the IEEE*, 108(10): 1785–1824, 2020. doi: 10.1109/JPROC.2020.2999034.

Travis Mortlock, Deevakar Muthirayan, Shanglin Y. Yu, Pramod P. Khargonekar, and Mohammad A. Al Faruque. Graph learning for cognitive digital twins in manufacturing systems. *IEEE Transactions on Emerging Topics in Computing*, 10(1):34–45, 2021. doi: 10.1109/TETC.2021.3132251.

Ziyi Ni, Minglun Han, Feilong Chen, Linghui Meng, Jing Shi, Pin Lv, and Bo Xu. Vilas: Exploring the effects of vision and language context in automatic speech recognition. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 11366–11370. IEEE, 2024.

Haochen Qi, Zhiwen Hu, Zhongliang Yang, Jian Zhang, Jie Jayne Wu, Cheng Cheng, Chunchang Wang, and Lei Zheng. Capacitive aptasensor coupled with microfluidic enrichment for real-time detection of trace sars-cov-2 nucleocapsid protein. *Analytical chemistry*, 94(6):2812–2819, 2022.

Daozheng Qu and Yanfei Ma. Magnet-bn: markov-guided bayesian neural networks for calibrated long-horizon sequence forecasting and community tracking. *Mathematics*, 13(17):2740, 2025.

Raúl Sánchez-Julián, Ignacio Lacalle, Rafael Vañó, Fernando Boronat, and Carlos E. Palau. Self-* capabilities of cloud-edge nodes: A research review. *Sensors*, 23(6):2931, 2023. doi: 10.3390/s23062931.

Concetta Semeraro, Mario Lezoche, Hervé Panetto, and Michele Dassisti. Digital twin paradigm: A systematic literature review. *Computers in Industry*, 130:103469, 2021. doi: 10.1016/j.compind.2021.103469.

Samad M. E. Sepasgozar. Differentiating digital twin from digital shadow: Elucidating a paradigm shift to expedite a smart, sustainable built environment. *Buildings*, 11(4):151, 2021. doi: 10.3390/buildings11040151.

Diego Souza, Thais Webber, and Elizabeth Wanner. Towards federated, autonomous and cognitive digital twins with darling. In *Proceedings of the 2025 Annual System of Systems Engineering Conference (SoSE)*, pp. 1–6. IEEE, 2025. doi: 10.1109/SoSE66311.2025.11083783.

Georgia Stavropoulou, Konstantinos Tsitseklis, Lydia Mavraidi, Kuo-I Chang, Anastasios Zafeiropoulos, Vasileios Karyotis, and Symeon Papavassiliou. Digital twin meets knowledge graph for intelligent manufacturing processes. *Sensors*, 24(8):2618, 2024. doi: 10.3390/s24082618.

Fei Tao, Qiang Qi, Lingling Wang, and Andrew Y. C. Nee. Digital twins and cyber–physical systems toward smart manufacturing and industry 4.0: Correlation and comparison. *Engineering*, 5(4): 653–661, 2019. doi: 10.1016/j.eng.2019.01.014.

Huanhuan Wang, Xiao Zhang, Youbing Xia, and Xiang Wu. An intelligent blockchain-based access control framework with federated learning for genome-wide association studies. *Computer Standards & Interfaces*, 84:103694, 2023.

Xiang Wu, Huanhuan Wang, Wei Tan, Dashun Wei, and Minyu Shi. Dynamic allocation strategy of vm resources with fuzzy transfer learning method. *Peer-to-Peer Networking and Applications*, 13 (6):2201–2213, 2020.

Xiang Wu, Huanhuan Wang, Yongting Zhang, Baowen Zou, and Huaqing Hong. A tutorial-generating method for autonomous online learning. *IEEE Transactions on Learning Technologies*, 17:1532–1541, 2024.

Ibrahim Yitmen, Seyed Alizadehsalehi, Ilgin Akıner, and Mehmet E. Akıner. An adapted model of cognitive digital twins for building lifecycle management. *Applied Sciences*, 11(9):4276, 2021. doi: 10.3390/app11094276.

APPENDIX

A RELATED WORK

A.1 TRADITIONAL DIGITAL TWIN FRAMEWORKS

Digital twin technology has evolved from product lifecycle management toward data-driven cyber-physical synchronization (Barricelli et al., 2019; Fuller et al., 2020; Qi et al., 2022). Foundational surveys consolidate definitions, capabilities, and application patterns, emphasizing the transition from static mirroring to dynamic behavior modeling and closed-loop optimization (Jones et al., 2020; Semeraro et al., 2021; Cao et al., 2025). The distinction between digital twins and digital shadows further clarifies requirements for bidirectional interactivity and autonomous behavior (Sepasgozar, 2021).

In the IoT context, enabling technologies—streaming ingestion, scalable storage, and analytics pipelines—underpin industrial platforms for monitoring and predictive maintenance (Minerva et al., 2020; Liu et al., 2024). While these systems demonstrate value in asset health management, they often rely on pre-programmed rules and task-specific models, limiting adaptivity and generalized reasoning. Domain deployments in buildings and construction highlight opportunities for richer semantics and lifecycle integration (Yitmen et al., 2021; El Mokhtari et al., 2022; Ni et al., 2024). Positioning digital twins within Industry 4.0 and cyber-physical systems underscores the need for context awareness, resilience, and real-time decision support (Tao et al., 2019; Jeremiah et al., 2024).

A.2 COGNITIVE COMPUTING AND AI INTEGRATION

The concept of cognitive digital twins advances twins from passive data dashboards to autonomous, learning systems that perceive, reason, and act (Ali et al., 2021; Lu et al., 2022). Conceptual frameworks emphasize knowledge-centric modeling and resilience, integrating symbolic knowledge with data-driven models for robust decision-making (Eirinakis et al., 2022; Guo et al., 2023). Prior studies on dynamic resource allocation with transfer learning (Wu et al., 2020) and tutorial generation for autonomous learning (Wu et al., 2024) investigated adaptive mechanisms for system management under changing operational conditions.

Knowledge graphs have emerged as a unifying substrate to encode entities, relations, constraints, and provenance, with recent work exploring automatic KG construction from sensor streams and LLM-assisted pipelines (Stavropoulou et al., 2024; Mandal & O'Connor, 2024b;a). Deep reinforcement learning and multi-agent coordination support adaptive task allocation in dynamic settings (Kim et al., 2020; Lee et al., 2022; Li et al., 2024). Graph learning within twins enables relational pattern discovery across assets, improving diagnosis and prediction (Mortlock et al., 2021). Representative instantiations include cognitive twins for smart manufacturing and unmanned maintenance (Guo et al., 2023; Lv et al., 2023; Al Faruque et al., 2021).

A.3 FEDERATED LEARNING AND DISTRIBUTED INTELLIGENCE

Federated learning (FL) enables privacy-preserving collaboration across distributed data silos, allowing multiple parties to jointly train models without sharing raw data (Kairouz et al., 2021; Li et al., 2020). Core advances address optimization under non-IID data, communication efficiency, and security guarantees. Prior work on blockchain-based access control with federated learning (Wang et al., 2023) and Markov-guided Bayesian networks for long-horizon forecasting (Qu & Ma, 2025) provided foundations for privacy-preserving distributed learning.

In industrial contexts, FL research has evolved to investigate encrypted, distributed training protocols accounting for edge device constraints (Kuo & Yang, 2024; Dritsas & Trigka, 2025). For cognitive digital twins, FL must interoperate with knowledge representation and real-time control loops. Recent work explores specialized privacy-preserving protocols and architectural designs for autonomous, federated twins capable of collaborative learning and coordinated action (Mandal, 2024; Souza et al., 2025; Sánchez-Julián et al., 2023).

B Implementation Details

B.1 Network Architecture

The CPE employs modality-specific encoders: a 1D-CNN with residual connections for sensor data, a ResNet-18 backbone for images, a 2-layer LSTM for sequential data, and a DistilBERT encoder for text. All encoders project to a shared 256-dimensional latent space. The attention fusion module uses 8 heads with $d_k = 32$.

The HDMM reactive layer uses a 3-layer MLP (256-128-64) with ReLU activations. The deliberative layer employs a 2-layer GRU for state prediction and a differentiable MPC solver with horizon $H = 10$. The reflective layer uses MAML with inner learning rate 0.01 and outer learning rate 0.001.

B.2 Training Configuration

All models are trained using the AdamW optimizer with learning rate 10^{-4} and weight decay 10^{-5}. Batch size is 64 for sensor/sequence data and 32 for image data. Training uses mixed-precision (FP16) on NVIDIA A100 GPUs. Federated learning uses 10 local epochs per round with 100 communication rounds.

B.3 Computational Resources

Experiments were conducted on a cluster with 8$\times$ NVIDIA A100 (80GB) GPUs and 128 CPU cores. Full CDT training requires approximately 24 hours for the manufacturing testbed, 8 hours for building automation, and 16 hours for supply chain.