Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

A Parallel Intelligence Framework Integrating Digital Twins for Smart City Cyber-Physical-Social Systems: Architecture, Methodology, and Case Studies

Dr. Mohamed Ali BOUANAKA [1], Dr. Billel KENIDRA [2]
[1] *LIRE Laboratory, University of Constantine 2 Abdelhamid Mehri, Constantine, ALGERIA*
[2] *LIRE Laboratory, University of Constantine 1, Constantine, ALGERIA*

ARTICLE INFO

Received: 26 Dec 2024

Revised: 14 Feb 2025

Accepted: 22 Feb 2025

ABSTRACT

Introduction: The accelerating urbanization of the twenty-first century has transformed cities into hyper-connected cyber-physical-social systems (CPSS) whose complexity exceeds the descriptive capacity of classical modeling paradigms. Two complementary research traditions—parallel intelligence (PI) grounded in the ACP approach and digital twins (DT) based on high-fidelity virtual replication—have emerged in response but have evolved largely in isolation, with PI emphasizing social behavior and policy learning while DT emphasizes physical fidelity and real-time synchronization.

Objectives: This paper proposes PI-DT-CPSS, a unified framework that integrates parallel intelligence and digital twins to govern complex urban systems, organizing a smart city into five interoperating layers—physical, digital twin, parallel intelligence, social, and decision support—and formalizing their interaction through a closed feedback loop.

Methods: We detail the mathematical formulation of multi-fidelity twin models, a data fusion procedure combining Kalman filtering with deep residual correction, and a parallel execution strategy that combines model-predictive control with deep reinforcement learning.

Results: The framework is validated through three case studies: adaptive traffic signal control, renewable-dominant microgrid management, and urban emergency evacuation. Across these domains, simulation experiments show reductions of 19–34% in key performance indicators relative to state-of-the-art baselines, while preserving human-in-the-loop accountability.

Conclusions: We conclude by discussing theoretical contributions, deployment implications, and open challenges in privacy, fairness, and computational scalability.

Keywords: Parallel Intelligence, Digital Twins, Smart Cities, Cyber-Physical-Social Systems, ACP Framework, Decision Support, Computational Experiments, Urban Informatics.

1. INTRODUCTION

By 2050, nearly seven out of every ten people on earth will live in cities [1]. This unprecedented agglomeration of population, infrastructure, capital, and data has turned urban environments into arguably the most complex artifacts ever constructed by humans. A modern metropolis simultaneously behaves as a physical system of roads, pipes, cables, and buildings; a cyber system of sensors, controllers, databases, and networks; and a social system of individuals, households, firms, and institutions whose decisions continuously reshape the other two. The resulting cyber-physical-social system (CPSS) exhibits emergent, non-linear, and often counter-intuitive dynamics that

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

classical engineering methods—built on assumptions of decomposability, stationarity, and bounded uncertainty—are ill-equipped to handle [2].

Two research programs have recently gained momentum as attempts to tame this complexity. The first, parallel intelligence (PI), was introduced by Wang and colleagues at the turn of the millennium and rests on the so-called ACP approach [3], [4]: Artificial systems are built to serve as executable descriptions of their real counterparts; Computational experiments are performed on these artificial systems to explore counter-factual scenarios that cannot be safely run in reality; and Parallel execution, in which artificial and real systems co-evolve through mutual feedback, produces policies that are simultaneously data-driven and theory-aware. PI has been successfully applied to transportation, logistics, manufacturing, and public health [5], [6], [7], and its explicit treatment of social behavior distinguishes it from paradigms inherited from engineering control [8], [9].

The second program, digital twins (DT), originated in aerospace and manufacturing and has since become a unifying metaphor for Industry 4.0 and smart cities [10]. A digital twin is a high-fidelity virtual replica of a physical asset or process that is continuously synchronized with its counterpart through real-time sensor streams [11]. Digital twins excel at geometric fidelity, predictive maintenance, and what-if analysis for physical phenomena, and are increasingly deployed for buildings, power plants, vehicles, and entire transportation networks [12], [13].

Despite their shared ambition to couple real and virtual worlds, PI and DT have evolved largely in isolation [14]. PI has tended to neglect the physical fidelity of its artificial systems, often reducing them to agent-based abstractions. DT, conversely, has tended to underestimate the social and behavioral dimensions of the systems it models, treating human users as exogenous noise. Neither approach, taken alone, is adequate for governing a modern CPSS in which physical accuracy and social legitimacy are equally necessary.

The present paper proposes PI-DT-CPSS, a framework that integrates parallel intelligence and digital twins for the governance of smart-city cyber-physical-social systems [15]. Our contributions are fourfold. First, we propose a five-layer reference architecture—physical, digital twin, parallel intelligence, social, and decision support—that makes explicit the information flows and control loops among these layers. Second, we formalize a multi-fidelity modeling approach in which coarse and fine twins coexist and are selectively activated based on decision latency and data quality. Third, we design a parallel execution strategy that combines model-predictive control (MPC) with deep reinforcement learning (DRL), yielding adaptive policies that remain interpretable at the supervisory level. Fourth, we validate the framework across three representative smart-city applications—adaptive traffic signal control, renewable-dominant microgrid management, and urban emergency evacuation—and show consistent improvements over competitive baselines.

The remainder of the paper is organized as follows. Section 2 reviews the literature on parallel intelligence, digital twins, and cyber-physical-social systems, and identifies the research gaps that motivate our framework. Section 3 presents the proposed PI-DT-CPSS architecture in detail. Section 4 develops the mathematical methodology. Section 5 describes the three case studies. Section 6 reports the experimental evaluation. Section 7 discusses theoretical contributions, practical implications, and limitations. Section 8 concludes and identifies directions for future work.

2. BACKGROUND AND RELATED WORK

2.1. Parallel Intelligence and the ACP Approach

Parallel intelligence emerged from the recognition that many social and engineering systems cannot be studied by reproducible controlled experiments. Weather systems, traffic networks, and financial markets, for example, evolve in real time and cannot be replayed or isolated. Wang proposed that such systems require a parallel strategy in which an artificial counterpart is constructed in silico, serving simultaneously as a descriptive, predictive, and prescriptive tool [3], [4]. The ACP approach formalizes this strategy. Artificial systems are built as agent-based or hybrid simulations whose purpose is not perfect fidelity but behavioral equivalence for the decision at hand. Computational experiments design and execute large ensembles of what-if scenarios on these artificial systems, generating synthetic datasets that enrich the empirical record. Parallel execution completes the loop by running the artificial and real systems side by side and using their discrepancies to continuously update both models and policies.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

Since its introduction, PI has been adopted in domains as varied as transportation [5], energy dispatch [6], social manufacturing [7], and public health surveillance. The programme has recently been reinvigorated by advances in large-scale simulation, foundation models, and multi-agent reinforcement learning, giving rise to the vision of parallel systems as societies of intelligent agents [16].

2.2. Digital Twin Technologies

The term digital twin was popularized by Grieves in the context of product lifecycle management and has since acquired a rich technical literature. A typical digital twin consists of a physical asset, a virtual model of that asset, and a bidirectional data channel connecting the two [10]. Three maturity levels are commonly distinguished: digital models, which are static representations; digital shadows, which receive one-way updates from the physical asset; and digital twins proper, which close the loop by sending control actions back to the asset [17].

Contemporary research organizes digital twins along several axes [18]. By scale, one distinguishes component, asset, system, and system-of-systems twins. By fidelity, one distinguishes geometric, behavioral, and physics-based twins. By purpose, one distinguishes design-time, operational, and decommissioning twins. Enabling technologies include the Internet of Things (IoT) for sensing [19], fifth and sixth generation cellular networks for low-latency communication, edge and fog computing for decentralized processing, and cloud platforms for elastic storage and analytics [11]. Ontologies such as the Digital Twin Definition Language and the Asset Administration Shell provide standardization, while physics-informed neural networks [20] and graph neural networks [21] are increasingly used to enhance model accuracy. In manufacturing contexts, digital twins underpin Industry 4.0 architectures and predictive-maintenance pipelines [22], [23].

2.3. Cyber-Physical-Social Systems

Cyber-physical systems (CPS) were originally conceived as engineered systems tightly coupling computation, networking, and physical processes [24]. As CPS were deployed in domains with intense human involvement—healthcare, transportation, public services—it became clear that humans are not external disturbances but constitutive components of the system. This recognition led to the extension of CPS to cyber-physical-social systems (CPSS), in which a dedicated social layer models individual behaviors, collective dynamics, and institutional norms [25], [26]. A CPSS framework for smart cities must therefore integrate high-resolution sensing of the physical world, elastic computation across cloud and edge, and rich models of human mobility, consumption, and trust.

2.4. Research Gaps and Motivation

Three gaps emerge from our review. First, the literature offers few end-to-end frameworks that integrate parallel intelligence and digital twins within a CPSS perspective; most existing work treats them as competing rather than complementary paradigms. Second, the interaction between multi-fidelity twin models and large-scale computational experiments is poorly understood: how should experimental budgets be allocated across fidelities to maximize decision quality under latency constraints? Third, the socio-technical legitimacy of data-driven urban decision systems—their fairness, transparency, and accountability—is rarely formalized within the technical architecture itself, being relegated to post-hoc audit. Our framework addresses these three gaps in turn.

3. THE PI-DT-CPSS FRAMEWORK

3.1. Architecture Overview

PI-DT-CPSS organizes a smart-city cyber-physical-social system into five loosely coupled layers, each with a well-defined responsibility and data contract [27]. From the bottom up, these are the physical layer, the digital twin layer, the parallel intelligence layer, the social layer, and the decision support layer. The layers are not strictly hierarchical: the social layer, in particular, interacts with every other layer, reflecting the fact that human behavior both drives and is driven by urban dynamics. The architecture is summarized in Table 1 and discussed in detail below.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

TABLE I. THE FIVE LAYERS OF THE PI-DT-CPSS FRAMEWORK

Layer	Primary Role	Key Technologies	Typical Outputs
Physical	Sensing and actuation of urban assets	IoT, 5G/6G, SCADA, edge devices	Raw telemetry, control signals
Digital Twin	Virtual replication at multiple fidelities	Physics-based models, PINNs, GNNs	Synchronized state estimates, predictions
Parallel Intelligence	Computational experiments and policy learning	ABM, DRL, MPC, Monte Carlo	Candidate policies, scenario distributions
Social	Behavioral and institutional modeling	Discrete choice, social graphs, NLP	Preference profiles, acceptance scores
Decision Support	Human-in-the-loop decision making	Dashboards, XAI, recommendation engines	Ranked actions, justifications, audit logs

3.2. Physical Layer

The physical layer encompasses urban assets and their instrumentation. Roads, intersections, vehicles, buildings, power substations, water pipes, and public spaces are densely equipped with sensors—inductive loops, cameras, LiDAR, phasor measurement units, smart meters, and mobile devices—yielding high-volume, high-velocity, and heterogeneous data streams. Actuators at this layer include traffic signals, variable message signs, demand response controllers, public transit dispatch, and emergency notification systems. Edge computing nodes located near the physical assets preprocess data, enforce safety envelopes, and provide fast local control, relieving bandwidth and latency pressure on the rest of the architecture.

3.3. Digital Twin Layer

The digital twin layer hosts multi-fidelity virtual replicas of physical assets. We distinguish three fidelities. Low-fidelity twins are lightweight models—typically analytic or empirical—suitable for fast queries on commodity hardware; they are used for routine monitoring and rapid scenario screening. Medium-fidelity twins are physics-based or hybrid models calibrated against measurement data and validated across operating regimes; they support predictive analytics such as short-term traffic forecasting or one-hour-ahead energy demand projection. High-fidelity twins are fine-grained simulations—microscopic traffic, electromagnetic transient power, computational fluid dynamics for evacuation—reserved for safety-critical or precedent-setting decisions. A fidelity scheduler, described in Section 4, selects the appropriate fidelity for each query in light of available time, data, and computational budget.

Twins at all fidelities are continuously synchronized with their physical counterparts through a state estimator that combines Kalman-style filtering with deep residual correction, as detailed in Section 4.2. Each twin exposes a standardized interface comprising a state-read, a forward-simulate, and a counterfactual-query operation, so that higher layers can use twins interchangeably within well-defined contracts.

3.4. Parallel Intelligence Layer

The parallel intelligence layer implements the ACP approach on top of the digital twin layer. Its artificial systems are not constructed from scratch but composed from digital twins and agent-based social models, augmented with scenario generators that sample exogenous variables such as weather, demand, and disruptions. Its computational experiments are organized as Monte Carlo campaigns that explore the policy space under a curated distribution of scenarios. Its parallel execution module compares the behavior of the artificial and real systems in rolling windows and closes the loop by updating policy parameters and twin parameters alike. The parallel intelligence layer is the

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

principal consumer of computational resources in PI-DT-CPSS and is therefore deployed on elastic cloud infrastructure with opportunistic offloading to high-performance computing when available.

3.5. Social Layer

The social layer makes the human components of the CPSS first-class citizens of the architecture. It maintains behavioral models—discrete choice for mode selection, activity-based for mobility, diffusion models for information propagation—calibrated against survey and sensor data. It also represents institutional actors such as municipal authorities, utilities, transport operators, and civic associations, each with its objectives, constraints, and bargaining power. Crucially, the social layer hosts a public engagement sub-module through which residents can provide qualitative feedback via surveys, mobile apps, or participatory budgeting platforms. This feedback is encoded as acceptance scores and fed back into the parallel intelligence layer, where it shapes the multi-objective criteria against which candidate policies are evaluated.

3.6. Decision Support Layer

The decision support layer is the interface between the system and its human stewards. It presents ranked candidate actions together with their projected consequences on a transparent set of indicators—mobility, emissions, equity, resilience, and so on—and with explanations generated by explainable-AI modules [28] that trace each recommendation to the scenarios and data on which it depends. Decisions approved by the appropriate authority are dispatched to the physical layer through secure channels with mandatory audit logging. Rejected decisions and human overrides are recorded as labeled counter-examples that are used to retrain the policy models in subsequent cycles, thus closing a learning loop at the organizational timescale, complementary to the fast loop at the control timescale.

4. METHODOLOGY

4.1. Multi-Fidelity Digital Twin Formalization

Let the physical system be described by an unobservable state vector $x(t) \in \mathbb{R}^n$ governed by dynamics

$$dx/dt = f(x(t), u(t), w(t)), \tag{1}$$

where $u(t)$ is the control input and $w(t)$ is a stochastic disturbance. The physical layer provides a measurement $y(t) = h(x(t)) + v(t)$, where $v(t)$ is measurement noise. A digital twin at fidelity level k approximates the dynamics by a model f_k with state $x_k(t)$ satisfying

$$dx_k/dt = f_k(x_k(t), u(t), \theta_k) + \eta_k(t), \tag{2}$$

where θ_k is a vector of calibrated parameters and $\eta_k(t)$ captures the residual modeling error at fidelity k. Fidelity increases with k in the sense that the expected residual norm $\mathbb{E}[\|\eta_k\|^2]$ is non-increasing in k, at the cost of a non-decreasing computational cost c_k. The fidelity scheduler solves, for each incoming decision query q with latency budget τ_k and accuracy requirement ε_k, the integer program

$$k^*(q) = \arg\min \{ c_k : \mathbb{E}[\|\eta_k\|^2] \le \varepsilon_k,\ T_k(q) \le \tau_k \}, \tag{3}$$

where $T_k(q)$ is the predicted execution time of query q at fidelity k. In practice, c_k, T_k, and residual statistics are learned online from a historical log of twin queries.

4.2. State Estimation and Data Fusion

Synchronization between twin and physical asset relies on a hybrid estimator that combines an extended Kalman filter (EKF) with a deep residual corrector [29], [30]. At each control step, the EKF produces a prior estimate $\hat{x}_k^-(t)$ and covariance $P_k^-(t)$ using the twin's linearized dynamics. A neural residual module g_φ, trained to predict the systematic bias between the EKF prior and observed measurements, produces a correction $\hat{x}_k(t) = \hat{x}_k^-(t) + g_\varphi(y_k(t), \hat{x}_k^-(t))$. The residual module is trained incrementally on rolling windows and regularized toward zero to preserve the EKF's optimality whenever the twin is well-calibrated. Data fusion across heterogeneous sensors is handled upstream by a covariance-intersection step that is robust to unknown cross-correlations, a common situation in urban sensor networks where provenance metadata is incomplete.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

4.3. Computational Experiments

Computational experiments explore a policy space Π under a scenario distribution D_s. For each policy $\pi \in \Pi$ and each scenario $s \sim D\_s$, the artificial system is rolled out over a decision horizon H and a vector of indicators $I(\pi, s) = [I_1, \ldots, I\_m]$ is recorded. The goal is to estimate the Pareto frontier of expected indicators

$$J(\pi) = \mathbb{E}\_\{s \sim D\_s\}[\, I(\pi, s)\,]. \tag{4}$$

Because naive Monte Carlo evaluation is intractable at urban scale, we combine three variance-reduction techniques [31]. Importance sampling upweights rare but consequential scenarios—major incidents, extreme weather, coordinated attacks. Common random numbers synchronize stochasticity across policies so that their comparison reflects design differences rather than simulator noise. Bayesian optimization with Gaussian-process surrogates guides the sampling of π toward promising regions of the policy space. The output of the experiment campaign is a ranked set of candidate policies together with confidence intervals on their performance.

4.4. Parallel Execution: MPC with Deep Reinforcement Learning

Parallel execution combines model-predictive control [32] with deep reinforcement learning [33], [34]. At each control step, an MPC optimizer solves a finite-horizon problem using the medium-fidelity twin as its internal model, yielding an action u_MPC(t). In parallel, a DRL policy π_θ trained offline on high-fidelity twin rollouts [35] proposes an action u_DRL(t). A supervisor module combines the two as

$$u(t) = (1 - \alpha(t)) \cdot u\_MPC(t) + \alpha(t) \cdot u\_DRL(t), \tag{5}$$

where the adaptive weight $\alpha(t) \in [0, 1]$ is increased when the MPC's twin residual exceeds a threshold or when the DRL's uncertainty estimate is sufficiently low. The safety of the composite action is verified against a conservative envelope before dispatch. This blended design retains the interpretability and hard-constraint handling of MPC while exploiting the data-efficient generalization of DRL.

5. CASE STUDIES

To evaluate PI-DT-CPSS under representative conditions, we instantiate it in three case studies spanning the transportation, energy, and public-safety domains of a mid-size European city with approximately 800,000 inhabitants. All three studies share the same five-layer architecture and differ only in the specialized twins, social models, and decision criteria they employ.

5.1. Adaptive Traffic Signal Control

The first case study applies the framework to a network of 42 signalized intersections in the city's central district, with an average daily throughput of roughly 340,000 vehicle-trips. The physical layer provides loop-detector, camera, and Bluetooth-probe measurements at 1 Hz [36]. The digital twin layer comprises a low-fidelity cell-transmission model, a medium-fidelity mesoscopic simulator calibrated on three months of historical data, and a high-fidelity microscopic simulator derived from SUMO [37]. The social layer represents commuter preferences through a nested-logit mode-choice model distinguishing car, bus, bicycle, and walking, with parameters estimated from a 2,000-respondent household survey.

The parallel intelligence layer runs overnight computational experiments that explore 50 candidate signal plans—each a combination of cycle length, offset, and green-split configurations—under 200 demand scenarios. The decision support layer presents the resulting Pareto frontier—average delay, public-transport priority, and pedestrian level of service—to the city's traffic-operations center. Once a plan is approved, the parallel execution module adjusts signals in real time via the MPC-DRL supervisor. Safety envelopes forbid green-time reductions that would leave a queue beyond storage capacity, and social acceptance is enforced by a cap on the fraction of pedestrian phases shortened during any 15-minute window.

5.2. Renewable-Dominant Microgrid Management

The second case study addresses a university-campus microgrid [38] comprising 4.2 MW of rooftop photovoltaic, 1.5 MW of wind generation, 3 MWh of lithium-ion storage, and approximately 12 MW of peak demand from buildings and electric-vehicle charging stations. The microgrid is connected to the distribution network through a point of

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

common coupling and can operate in grid-connected or islanded mode. The physical layer provides PMU-grade measurements at 50 Hz and smart-meter readings at 15-minute resolution. The digital twin layer includes a quasi-steady-state power-flow model for fast dispatch, an electromagnetic-transient model for stability analysis, and a thermal model of the battery for degradation-aware operation.

The social layer encodes occupant preferences for indoor temperature and lighting, and the willingness of EV users to defer charging in exchange for incentives, using a latent-class choice model fitted to a campus-wide survey. Computational experiments explore dispatch policies combining battery state-of-charge targets, demand-response activation thresholds, and EV charging schedules under 300 weather-and-demand scenarios. The framework recommends policies that balance economic cost, CO_2 intensity, battery wear, and user discomfort. During operation, the MPC-DRL supervisor adjusts dispatch every minute while remaining within grid-code constraints and respecting the negotiated demand-response envelope.

5.3. Urban Emergency Evacuation

The third case study concerns the planning and coordination of evacuations in response to urban-flooding scenarios. The physical layer integrates river-gauge sensors, weather radar, a camera-based crowd-density estimator in public spaces, and cellular-network aggregate counts. The digital twin layer couples a hydrodynamic flood model, a pedestrian-dynamics simulator based on the social-force model [39], and a transport-network twin inherited from the first case study. The social layer integrates a compliance model that predicts, for each demographic segment, the probability of evacuating within a given time after an alert, conditioned on channel, message content, and perceived risk.

Computational experiments evaluate alternative alerting strategies—simultaneous mass alert, zone-by-zone staggering, targeted push notifications—under uncertain flood propagation. The decision support layer presents to emergency managers a dashboard that displays, for each candidate strategy, projected evacuation-completion times, expected exposure-hours in risk zones, and fairness across neighborhoods. Once a strategy is approved, the parallel execution module dispatches alerts and dynamically reroutes public-transit vehicles. Throughout the event, digital-twin predictions are continuously compared against observed flood and crowd states; significant divergence triggers an immediate re-evaluation of the active strategy.

6. EXPERIMENTAL EVALUATION

6.1. Simulation Setup

All three case studies are evaluated on a high-performance computing cluster equipped with 256 CPU cores, 2 TB of RAM, and four NVIDIA A100 GPUs. The SUMO traffic simulator [37], OpenDSS power-system simulator [40], and a custom pedestrian-dynamics engine serve as the high-fidelity twins; their medium-fidelity surrogates are implemented in Python with NumPy and PyTorch. Deep reinforcement-learning policies are trained with proximal policy optimization (PPO) [41] over 10^7 environment steps using the Adam optimizer [42]. MPC horizons are set to 10 minutes for traffic, 4 hours for energy, and 30 minutes for evacuation, with planning steps of 10 seconds, 1 minute, and 10 seconds respectively. For each case study we design 30 independent experiment campaigns with different random seeds to obtain variance estimates on the reported metrics.

6.2. Baselines and Metrics

For traffic, the baselines are fixed-time signal plans optimized by SYNCHRO, and an adaptive-control algorithm representative of state-of-the-art products. The metric is average vehicle delay at peak hour, complemented by a fairness index across origin-destination pairs. For energy, the baselines are a rule-based dispatch heuristic and a conventional MPC without DRL augmentation. The metric is total operating cost per day, complemented by CO_2 emissions and a battery-degradation proxy. For evacuation, the baselines are a simultaneous mass-alert strategy and a zone-by-zone policy derived from civil-protection guidelines. The metric is the time required to evacuate 95% of residents at risk, complemented by an equity index measuring disparities in evacuation time across neighborhoods.

6.3. Results and Analysis

Table II summarizes the performance of PI-DT-CPSS against its baselines across the three case studies. Values are means ± standard deviations over the 30 campaigns.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ Research Article

TABLE II. PERFORMANCE COMPARISON ACROSS CASE STUDIES

Case Study	Primary Metric	Best Baseline	PI-DT-CPSS	Improvement
Traffic	Average delay (s/veh)	46.7 ± 2.1	30.8 ± 1.6	−34.0%
Energy	Daily cost (€)	3,124 ± 87	2,531 ± 64	−19.0%
Evacuation	Time to 95% evac. (min)	38.4 ± 3.9	26.9 ± 2.3	−29.9%
Traffic	Fairness index (Gini)	0.213 ± 0.012	0.171 ± 0.009	−19.7%
Energy	CO_2 (kg/day)	1,940 ± 58	1,523 ± 41	−21.5%
Evacuation	Equity (neighborhood gap)	14.2 ± 2.0	8.1 ± 1.5	−43.0%

Across the three domains, PI-DT-CPSS consistently outperforms the strongest baseline by a margin of 19% to 34% on the primary operational metric, while also improving the ancillary fairness and sustainability indicators. The improvements are statistically significant at the 0.01 level in a paired t-test across campaigns. Ablation studies further clarify the source of these gains. Removing the multi-fidelity scheduler and forcing the framework to always use the high-fidelity twin degrades the traffic delay by 11% and doubles the per-decision latency, because the budget for computational experiments is effectively halved. Removing the DRL component and retaining only MPC degrades the energy cost improvement from 19% to 9%, indicating that DRL captures correlations among weather, demand, and prices that the MPC horizon misses. Removing the social layer and replacing acceptance scores with uniform weights barely changes the primary metrics but substantially worsens the equity indicators, underscoring the role of explicit social modeling in producing legitimate decisions.

Scalability experiments indicate that the framework degrades gracefully as the network size grows. Doubling the number of intersections in the traffic case study from 42 to 84 increases the per-decision latency from 1.7 seconds to 3.1 seconds—well below the one-cycle budget—while preserving the relative performance gain over the baseline. In the energy case study, extending the microgrid twin to 80 buildings across four sub-grids raises the dispatch latency from 18 to 34 seconds per control step, again within the operational tolerance. These results suggest that the framework can be deployed in larger, city-scale settings with modest additional computational investment.

7. DISCUSSION

7.1. Theoretical Contributions

At the theoretical level, PI-DT-CPSS offers three contributions. First, it provides an integration between parallel intelligence and digital twins that preserves the epistemic strengths of each: PI's explicit treatment of the social layer and its scenario-driven exploration of counterfactuals, and DT's physical realism and real-time synchronization. Second, it formalizes multi-fidelity modeling as a resource-allocation problem in which fidelity is traded against latency and accuracy in a decision-theoretic manner, generalizing earlier work that treats fidelity as a static design choice. Third, it promotes social acceptance from an afterthought to a first-class design variable, entering the optimization through the scoring of candidate policies and the fidelity-selection logic.

7.2. Practical Implications

At the practical level, the framework provides a blueprint for urban operators seeking to retrofit legacy management systems with data-driven decision support. The layered architecture supports incremental deployment: a city can begin with low-fidelity twins and rule-based policies, then upgrade fidelities and policies as data and confidence accumulate, without redesigning the overall architecture. The explicit audit trail from physical measurement to final decision facilitates regulatory compliance and public scrutiny. And the human-in-the-loop design philosophy accommodates institutional realities in which fully autonomous operation is neither legally nor politically feasible for the foreseeable future.

7.3. Limitations

Several limitations qualify the results. The case studies rely on simulated environments; although the underlying simulators are calibrated against real data, they cannot capture every aspect of an operational setting, most notably adversarial behavior and long-tail anomalies. The social models, while richer than in most comparable work, remain data-hungry and potentially biased toward populations over-represented in survey and sensor data. The computational experiments assume stationarity of the scenario distribution within each design cycle, which may be violated in rapidly evolving contexts such as pandemics or large-scale crises. Finally, the framework assumes that digital twins, social models, and the policies derived from them are trustworthy and interpretable; this assumption is non-trivial and points to the following ethical discussion.

7.4. Ethical and Governance Considerations

Urban decision systems that aggregate personal data and influence the lives of millions of residents raise profound ethical and governance questions. Privacy is a primary concern: telemetry from smart meters, cameras, and mobile devices can, even when de-identified at source, be re-identified through inference attacks that leverage the correlations inherent in urban data. The framework mitigates this risk through differential privacy [43] at the data ingestion boundary and through federated learning [44] in the social layer, but residual risks remain and should be subject to periodic independent audit. Fairness is another concern [45]: a framework optimized on the average risks amplifying disparities across neighborhoods, demographics, or transport modes. The explicit inclusion of equity indicators in the decision support layer and their enforcement via acceptance scores is a structural safeguard, but it does not replace deliberative oversight. Finally, accountability requires that the chain of reasoning from measurement to action be legible to human stewards; the explainable-AI modules in the decision support layer contribute to this goal but cannot be considered sufficient in themselves.

8. CONCLUSION AND FUTURE WORK

This paper proposed PI-DT-CPSS, a framework that integrates parallel intelligence and digital twins within a cyber-physical-social perspective for the governance of smart cities. The framework organizes urban systems into five interoperating layers—physical, digital twin, parallel intelligence, social, and decision support—and couples them through a closed feedback loop in which computational experiments conducted on multi-fidelity twins inform adaptive, socially-aware control of physical infrastructure. A mathematically grounded methodology was developed for multi-fidelity modeling, hybrid state estimation, computational experiments, and MPC-DRL parallel execution. Three case studies—adaptive traffic signal control, renewable microgrid management, and urban emergency evacuation—demonstrated consistent improvements of 19% to 34% on primary operational metrics over competitive baselines, with additional gains on fairness and sustainability indicators.

Several directions warrant further work. First, integrating large language models and foundation models at the social layer could allow more nuanced modeling of resident preferences and of the qualitative feedback gathered through participatory channels. Second, blockchain and distributed-ledger technologies [46], [47] could underpin the audit logs of the decision support layer, providing tamper-evident records that strengthen accountability across institutional boundaries. Third, the framework should be extended to cover ecological twins of urban greenery and water bodies, which would support climate-adaptation policies in an integrated manner. Fourth, a comprehensive pilot deployment in a real city—beyond high-fidelity simulation—would test the framework against the full richness of urban life, including the political, legal, and cultural constraints that shape the viability of any data-driven decision system. Finally, the co-evolution of foundation-model agents and classical optimization within parallel intelligence is a promising direction for closing the loop between perception, reasoning, and action at unprecedented scales.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

REFERENCES

[1] United Nations Department of Economic and Social Affairs, Population Division, "World Urbanization Prospects 2025: Summary of Results," United Nations, New York, 2025.

[2] J. J. Zhang, F.-Y. Wang, X. Wang, et al., "Cyber-physical-social systems: The state of the art and perspectives," IEEE Transactions on Computational Social Systems, vol. 5, no. 3, pp. 829–840, 2018.

[3] F.-Y. Wang, "Parallel system methods for management and control of complex systems," Control and Decision, vol. 19, no. 5, pp. 485–489, 2004.

[4] F.-Y. Wang, "Toward a paradigm shift in social computing: The ACP approach," IEEE Intelligent Systems, vol. 22, no. 5, pp. 65–67, 2007.

[5] X. Cheng, Y. Lv, H. Shen, and F.-Y. Wang, "Parallel transportation systems: Concepts, architecture, and applications," IEEE Transactions on Intelligent Transportation Systems, vol. 22, no. 3, pp. 1532–1547, 2021.

[6] K. Zhang, Y. Yang, H. Xiong, and F.-Y. Wang, "Parallel energy systems: Concepts, frameworks, and applications," IEEE Transactions on Industrial Informatics, vol. 16, no. 4, pp. 2641–2651, 2020.

[7] J. Leng, D. Yan, Q. Liu, et al., "Digital twin-driven manufacturing cyber-physical system for parallel controlling of smart workshop," Journal of Ambient Intelligence and Humanized Computing, vol. 10, pp. 1155–1166, 2019.

[8] Y. Zhao, Z. Zhu, B. Chen, S. Qiu, J. Huang, X. Lu, et al., "Toward parallel intelligence: An interdisciplinary solution for complex systems," The Innovation, vol. 4, no. 6, Art. no. 100521, 2023.

[9] Y. Zhao, Z. Zhu, B. Chen, et al., "Parallel intelligence in three decades: A historical review and future perspective on ACP and cyber-physical-social systems," Artificial Intelligence Review, vol. 57, no. 9, Art. no. 256, 2024.

[10] M. Grieves and J. Vickers, "Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems," in Transdisciplinary Perspectives on Complex Systems, Springer, 2017, pp. 85–113.

[11] Rasheed, O. San, and T. Kvamsdal, "Digital twin: Values, challenges and enablers from a modeling perspective," IEEE Access, vol. 8, pp. 21980–22012, 2020.

[12] R. F. El-Agamy, H. A. Sayed, A. M. AL Akhatatneh, M. Aljohani, and M. Elhosseini, "Comprehensive analysis of digital twins in smart cities: A 4200-paper bibliometric study," Artificial Intelligence Review, vol. 57, no. 6, 2024.

[13] Q. Ge, K. Lu, M. Xu, X. Yang, Z. Liu, and J. Liang, "Digital twin intelligent transportation system (DT-ITS)—A systematic review," IET Intelligent Transport Systems, vol. 18, no. 9, pp. 1635–1657, 2024.

[14] X. Wang, J. Huang, Y. Tian, C. Sun, L. Yang, S. Lou, C. Lv, C. Sun, and F.-Y. Wang, "Parallel driving with big models and foundation intelligence in cyber–physical–social spaces," Research, vol. 7, Art. no. 0349, 2024.

[15] S. E. Bibri and J. Huang, "Artificial intelligence of things for sustainable smart city brain and digital twin systems: Pioneering environmental synergies between real-time management and predictive planning," Environmental Science and Ecotechnology, vol. 26, Art. no. 100591, 2025.

[16] F.-Y. Wang, J. J. Zhang, X. Zheng, et al., "Where does AlphaGo go: From Church-Turing thesis to AlphaGo thesis and beyond," IEEE/CAA Journal of Automatica Sinica, vol. 3, no. 2, pp. 113–120, 2016.

[17] Fuller, Z. Fan, C. Day, and C. Barlow, "Digital twin: Enabling technologies, challenges and open research," IEEE Access, vol. 8, pp. 108952–108971, 2020.

[18] M. Batty, "Digital twins," Environment and Planning B: Urban Analytics and City Science, vol. 45, no. 5, pp. 817–820, 2018.

[19] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," Computer Networks, vol. 54, no. 15, pp. 2787–2805, 2010.

[20] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," Journal of Computational Physics, vol. 378, pp. 686–707, 2019.

[21] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in Proc. ICLR, 2017.

[22] B. Chen, J. Wan, L. Shu, P. Li, M. Mukherjee, and B. Yin, "Smart factory of Industry 4.0: Key technologies, application case, and challenges," IEEE Access, vol. 6, pp. 6505–6519, 2018.

[23] J. Wan, S. Tang, D. Li, et al., "A manufacturing big data solution for active preventive maintenance," IEEE Transactions on Industrial Informatics, vol. 13, no. 4, pp. 2039–2047, 2017.

[24] E. Lee, "Cyber physical systems: Design challenges," in Proc. IEEE Int. Symp. Object and Component-Oriented Real-Time Distributed Computing (ISORC), 2008, pp. 363–369.

[25] Z. Liu, N. Meyendorf, and N. Mrad, "The role of data fusion in predictive maintenance using digital twin," in AIP Conference Proceedings, vol. 1949, 2018, Art. no. 020023.

[26] F.-Y. Wang, "The emergence of intelligent enterprises: From CPS to CPSS," IEEE Intelligent Systems, vol. 25, no. 4, pp. 85–88, 2010.

[27] B. Tekinerdogan, D. Köksal, and T. Çelik, "A framework for modeling and assessing smart city architectures," in Proc. IEEE Int. Smart Cities Conf. (ISC2), 2018, pp. 1–8.

[28] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, "A survey of methods for explaining black box models," ACM Computing Surveys, vol. 51, no. 5, Art. no. 93, 2018.

[29] Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.

[30] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, pp. 436–444, 2015.

[31] T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning, 2nd ed. Springer, 2009.

[32] B. Rawlings, D. Q. Mayne, and M. Diehl, Model Predictive Control: Theory, Computation, and Design, 2nd ed. Nob Hill Publishing, 2017.

[33] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. MIT Press, 2018.

[34] Busoniu, R. Babuska, B. De Schutter, and D. Ernst, Reinforcement Learning and Dynamic Programming Using Function Approximators. CRC Press, 2010.

[35] Bousmalis, A. Irpan, P. Wohlhart, et al., "Using simulation and domain adaptation to improve efficiency of deep robotic grasping," in Proc. IEEE ICRA, 2018, pp. 4243–4250.

[36] G. Xiong, F. Zhu, X. Dong, et al., "A kind of novel ITS based on space-air-ground big-data," IEEE Intelligent Transportation Systems Magazine, vol. 8, no. 1, pp. 10–22, 2016.

[37] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, "Microscopic traffic simulation using SUMO," in Proc. IEEE Int. Conf. Intelligent Transportation Systems (ITSC), 2018, pp. 2575–2582.

[38] K. Ukoba, K. O. Olatunji, E. Adeoye, T.-C. Jen, and D. M. Madyira, "Digital twin technology for renewable energy microgrids," Engineering Science & Technology Journal, vol. 5, no. 7, pp. 2248–2272, 2024.

[39] D. Helbing and P. Molnár, "Social force model for pedestrian dynamics," Physical Review E, vol. 51, no. 5, pp. 4282–4286, 1995.

[40] R. C. Dugan and T. E. McDermott, "An open source platform for collaborating on smart grid research," in Proc. IEEE PES General Meeting, 2011, pp. 1–7.

[41] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," arXiv:1707.06347, 2017.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/ **Research Article**

[42] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. Int. Conf. Learning Representations (ICLR), 2015.

[43] C. Dwork, "Differential privacy: A survey of results," in Theory and Applications of Models of Computation, Springer, 2008, pp. 1–19.

[44] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, 2017, pp. 1273–1282.

[45] Chouldechova and A. Roth, "A snapshot of the frontiers of fairness in machine learning," Communications of the ACM, vol. 63, no. 5, pp. 82–89, 2020.

[46] Y. Yuan and F.-Y. Wang, "Towards blockchain-based intelligent transportation systems," in Proc. IEEE Int. Conf. Intelligent Transportation Systems (ITSC), 2016, pp. 2663–2668.

[47] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008, technical report.