Digital Twins in Intelligent Transportation and Communication Systems: A Survey

Zifan Zhang, Dianwei Chen, Xianfeng Yang, *Member, IEEE,*
Shiwen Mao, *Fellow, IEEE,* Yuchen Liu, *Member, IEEE*

***Abstract*—Digital Twins (DTs) are reshaping Intelligent Transportation and Communication Systems (ITCS) by providing continuously synchronized, executable models of vehicles, infrastructure, and communication networks. This survey offers a unified taxonomy and explicitly defines six DT categories from design perspective: Single-Vehicle DT (SV-DT), a per-vehicle twin that mirrors dynamics, perception, planning, and control for inner-loop analysis of the vehicle; Connected Vehicle DT (CV-DT), a fleet or multi-vehicle twin that coordinates interactions, cooperation, and resource sharing under communications; Environment DT (ENV-DT), a twin of roads, intersections, signals, and demand that supports monitoring, forecasting, and policy evaluation at corridor and system scales; Network DT (NET-DT), a communication twin that emulates V2X and backhaul state to assess latency, reliability, and capacity and to guide roadside unit placement; Safety and Security DT (SAFE-DT), a cross-cutting twin for hazard analysis, incident detection, robustness, and cyber resilience; and finally, Twin of Twins (ToT), an orchestration layer that composes heterogeneous twins to deliver system-level services with consistency. This survey formalizes a comprehensive modular DT architecture in ITCS, including data paths, modeling fidelity, and system synchronization, and distinguishes DTs from simulators and digital shadows. Open challenges are also discussed along with future directions, with the goal of developing efficient and sustainable next-generation intelligent vehicular network systems.**

***Index Terms*—Digital twins, intelligent transportation and communication systems, vehicular networks**

I. Introduction

Intelligent transportation and communication systems (ITCS) have advanced from isolated traffic control and offline simulation to a tightly integrated cyber–physical infrastructure that couples connected and automated vehicles (CAVs), roadway sensors, roadside units (RSUs), and edge computing through Vehicle-to-Everything (V2X) communications [1], [2]. This evolution is driven by advances in multi-modal perception hardware such as Light Detection and Ranging (LiDAR), radar, and cameras, by precise mapping and localization, and by scalable data-driven pipelines deployed at the edge with multi-access edge computing (MEC) and in the cloud infrastructure [3]–[5]. In parallel, digital twin (DT) paradigm, originating in manufacturing and later generalized to cyber–physical systems, has matured into data-driven models that are continuously synchronized with physical assets and processes [6], [7]. In transportation settings, DTs provide machine-interpretable representations of vehicles, infrastructure, and communication networks that can be calibrated online, integrated with co-simulation tool chains, and coupled to physical ITCS in a closed loop, which supports analysis, prediction, and control of vehicles across multiple spatial and temporal scales [8], [9].

With advancements in machine learning (ML) and data-driven optimization methods [10], [11], DTs are able to deliver capabilities that conventional simulation and offline analytics cannot match in fidelity, timeliness, or operational relevance. By maintaining a continuously calibrated representation of the entire transportation system, DTs enable closed-loop decision making, systematic what-if and counterfactual evaluation, and safe rehearsal of safety-critical edge cases without exposing public roads to risk [12]–[15]. Heterogeneous data streams from onboard sensors and roadside infrastructure can be assimilated to perform state estimation and parameter identification, allowing DTs to track latent system variables and adapt their models as conditions evolve. This supports short-horizon trajectory and network forecasts at millisecond-to-second scales for vehicle control, as well as minute-to-hour horizons for traffic operations and planning [6], [7]. Practical benefits include predictive maintenance and asset management for fleets and RSUs, rapid prototyping and regression testing of autonomy or traffic control stacks, hardware-in-the-loop evaluation for perception and planning, and policy analysis for interventions such as dynamic lane allocation, variable speed limits, and pricing. For operators and service providers, DTs provide auditable interfaces for Quality of Service (QoS) and Quality of Experience (QoE), enabling proactive resource allocation, incident response, and resilience against faults and cyberattacks, while offering traceability for compliance and safety cases through reproducible logs and scenario catalogs [16]. Uncertainty quantification and explainability can be embedded to expose confidence in predictions and decisions, which is essential for risk management and certification [17]. At the system level, DTs facilitate multi-objective optimization that balances travel time, energy consumption, and emissions, and they support privacy-aware workflows by enabling synthetic data generation and evaluation without broad redistribution of raw sensing data [18], [19]. As a unifying substrate across the hierarchical networked system, DTs align single vehicle analytics with cooperative vehicle coordination, environment and infrastructure management, V2X communications, and safety assurance, creating an operational backbone that motivates a comprehensive survey of methods at these layers.

In this survey, a DT denotes a context-aware and persistently synchronized digital replica of a physical ITCS asset, process, and network that supports monitoring, prediction, and control, which is specified in ITU standards [20], [21]. It is also aligned with broader framework standards, such as ISO 23247 [22],

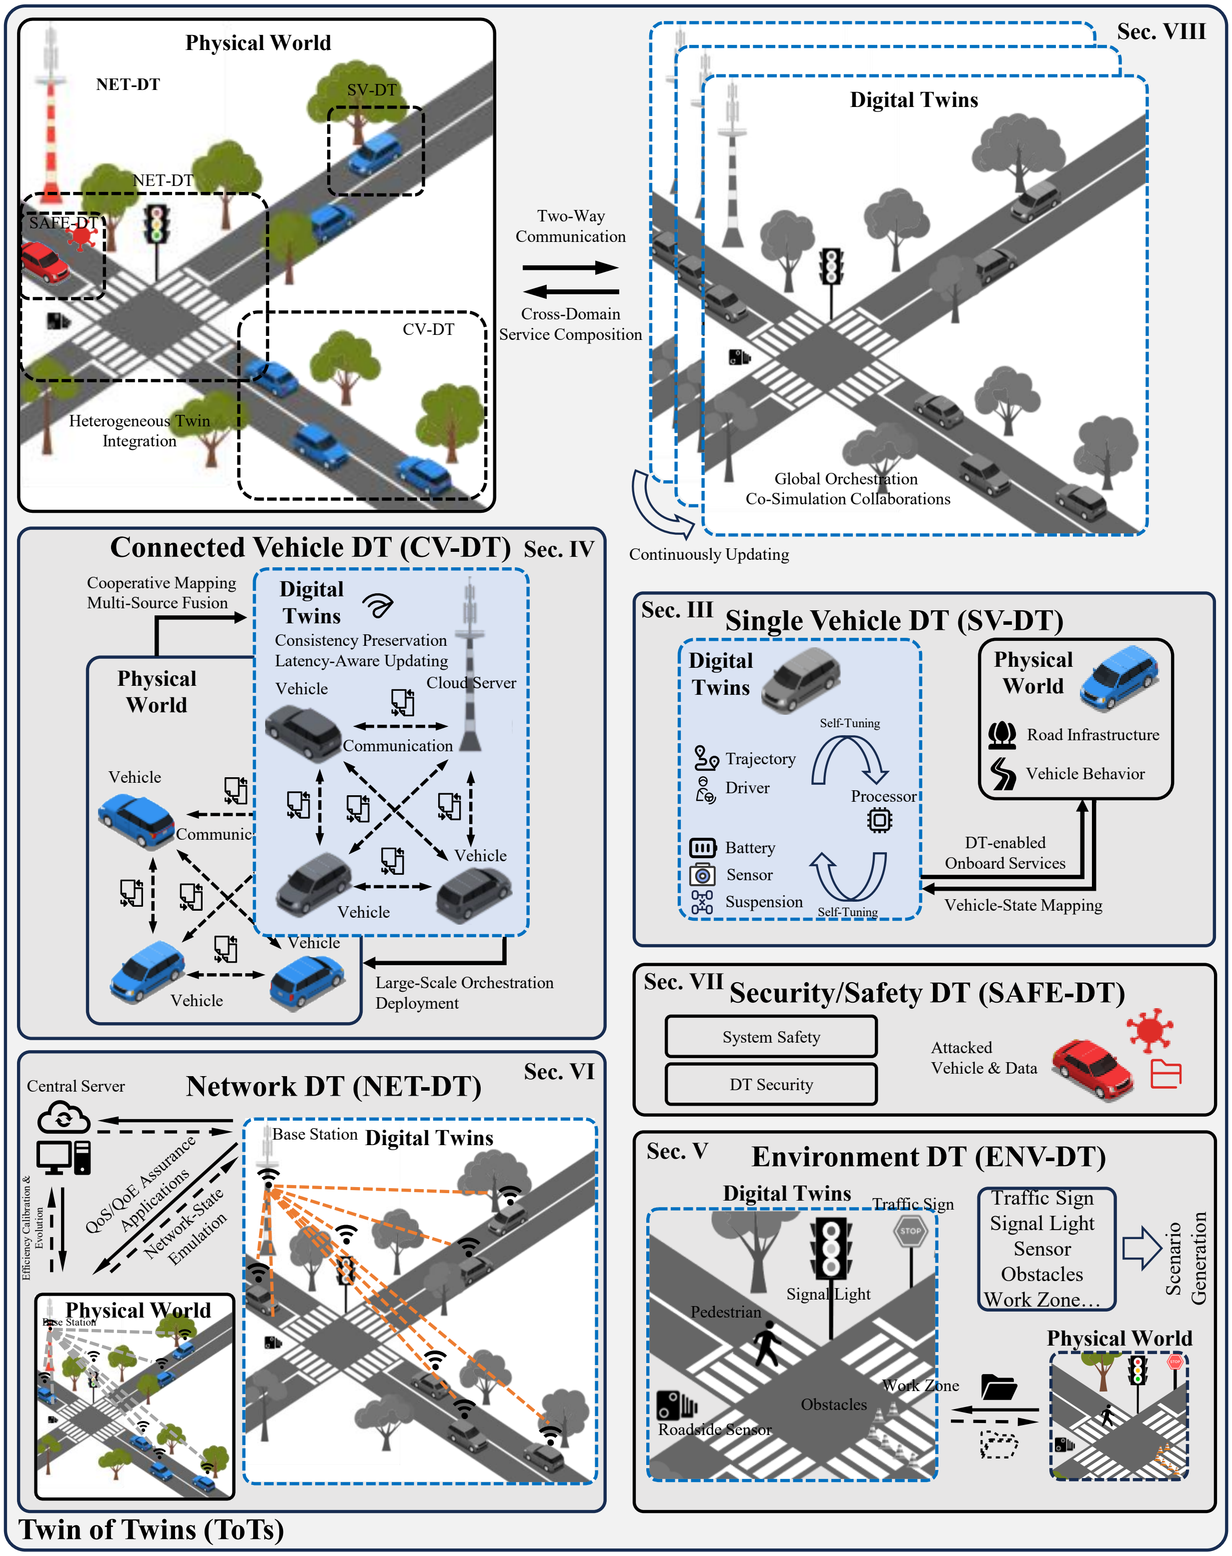

Fig. 1: Proposed taxonomy and architecture of DTs in ITCS, divided into six categories: (i) SV-DTs: per-vehicle twins coupling both digital and physical vehicle assets via self-tuning; (ii) CV-DTs: multi-vehicle twins with consistency preservation, on-demand communications, latency-aware updates, and large-scale orchestration; (iii) ENV-DTs: environment/infrastructure twins for traffic scenario generation and asset monitoring; (iv) NET-DTs: network twins for V2X with QoS/QoE assurance and network-wide emulation; (v) SAFE-DTs: safety and security twins spanning ITCS operational safety and adversarial resilience. (vi) All modules are integrated as ToTs: cross-domain composition, heterogeneous inter-twin connection, and global co-simulation. Bi-directional synchronization links the physical system and each DT module, and ToTs coordinates cross-twin services.

3GPP TR 22.842 [23], and ASAM OpenDRIVE [24]. A DT is distinguished from a conventional static simulator by continuous data assimilation, low latency, and fidelity-aware synchronization, and the ability to close the loop with decision or actuation when appropriate. This survey covers all components within ITCS, including vehicles, drivers, road and intersection infrastructure, traffic management processes, and communication networks, together with safety and security workflows and cross-domain orchestration. We specifically focus on studies that instantiate and maintain a DT using continuous telemetry data, co-simulation, physical deployment, data-driven communication and management.

Unlike general DTs that typically model isolated assets or processes, DTs for ITCS exhibit fundamentally different architectural, networking, and control characteristics. **The key distinctions** are summarized below.

- **Communication-Centric by Design.** In ITCS, the communication system is not auxiliary infrastructure but the foundational substrate that enables perception, coordination, control, and safety. Vehicular DTs rely on V2V, V2I, and V2X communication for real-time synchronization, state dissemination, and cooperative decision-making. Communication exists both as *intra-twin communication* (e.g., sensor-to-edge and vehicle-internal subsystems) and *inter-twin communication* (e.g., single-vehicle twin to network twin interactions). Unlike traditional industrial DTs where networking is merely a data pipe, ITCS-DTs embed networking as an integral and optimizable component of the twin architecture.
- **Hierarchical and Modular Twinning Architecture.** ITCS-DTs are inherently multi-layered and modular, forming a composable Twin-of-Twins structure. A complete system requires coordinated modeling of Single-Vehicle DT (SV-DT), Connected-Vehicle DT (CV-DT), Environment DT (ENV-DT), Network DT (NET-DT), and Safety DT (SAFE-DT). These layers exhibit strong hierarchical dependencies: vehicle-level decisions depend on environmental perception, which in turn depends on communication reliability and network state. This hierarchical coupling differentiates ITCS-DTs from conventional DTs that often focus on single-layer asset modeling without cross-layer orchestration.
- **Closed-Loop Co-Design of Sensing, Communication, and Control.** General DTs frequently emphasize monitoring and predictive maintenance. In contrast, ITCS-DTs must support real-time closed-loop optimization across sensing, communication, and control loops. For example, channel quality influences perception fidelity, latency impacts platoon stability, and congestion affects cooperative safety margins. Consequently, vehicular DTs require joint modeling and co-design across physical dynamics and network behavior, making cross-layer optimization a defining feature.
- **Ultra-Dynamic and Safety-Critical Operation.** Vehicular systems operate under high mobility, rapidly evolving topologies, and stringent latency and reliability constraints. ITCS-DTs must adapt at millisecond-level timescales while accounting for stochastic connectivity and dynamic channel conditions. Moreover, safety requirements impose probabilistic guarantees on collision risk, reliability, and fail-safe behavior. Such ultra-dynamic and safety-critical characteristics are significantly more demanding than those in relatively static DT deployments.
- **Metric-Coupled and Benchmark-Driven Evaluation.** In ITCS, each twin layer maps to measurable system-level metrics, such as trajectory deviation and energy efficiency for SV-DT, latency and reliability for NET-DT, and collision probability or risk bounds for SAFE-DT. These metrics are tightly coupled across layers, requiring joint evaluation frameworks rather than isolated performance indicators. This metric coupling highlights the need for unified taxonomies and benchmark alignment, which are often missing in prior DT surveys.

TABLE I: Survey positioning by scope and emphasis. A filled circle indicates substantive coverage, an open circle indicates partial coverage, and a dash indicates not covered.

Existing Surveys	**SV-DT**	**CV-DT**	**ENV-DT**	**NET-DT**	**SAFE-DT**	**ToT**
Bhatti et al. (2021) [25]	●	○	–	–	○	–
Deng et al. (2021) [26]	–	–	●	–	○	–
Hu et al. (2022) [27]	●	–	–	–	○	–
Wang et al. (2022) [9]	●	●	–	○	–	–
Hu et al. (2023) [8]	●	●	–	–	–	–
Irfan et al. (2024) [28]	–	○	○	○	●	–
Guo et al. (2024) [29]	–	●	–	●	–	–
Pan et al. (2025) [30]	–	○	–	●	–	–
This work	●	●	●	●	●	●

Table I summarizes positioning relative to previous surveys in coverage and depth. This survey, for the first time, integrates SV-DT, CV-DT, ENV-DT, NET-DT, SAFE-DT, and Twin-of-Twins (ToTs) layers, tying each to specific metrics and benchmarks to enable side-by-side comparison. The contribution of this survey can be summarized as follows:

- We propose a novel, unified taxonomy and literature review that jointly covers all ITCS-DT components from the design philosophy, enabling consistent comparison across layers of the hierarchical structure among functional DTs.
- We develop a standardized framework to review all stages of DTs, including construction, maintenance, deployment, and safety and security enhancement in ITCS.
- We perform a comprehensive analysis of open challenges and forward-looking directions that connect scalability, interoperability, verification and validation, data governance, and security to actionable research problems, with the goal of developing efficient and sustainable next-generation intelligent transportation and communication systems.

Figure 1 presents the overall architecture and taxonomy that guide the survey. Based on this structure, we organize DTs into six distinct categories by design, detailed in Sections III-A through VIII. We discuss each DT by its enabling technologies and tasks. The remainder of this survey is organized as follows: Sec. II introduces the fundamentals and background of ITCS and DT architectures, including their essential components. Secs III through VIII systematically present SV-DT, CV-DT, ENV-DT, NET-DT, SAFE-DT, and ToT, respectively, covering their specific construction, management, and decision-making processes. Sec. IX discusses open challenges, while Sec. X outlines corresponding future research directions. Finally, Sec. XI concludes the survey.

II. Fundamentals and Background

A. Preliminaries on DTs

Fundamentally, a DT is distinct from a static simulator and from a digital shadow. A simulator is an executable model that can generate trajectories under specified conditions but does not persistently assimilate live measurements during operation. A digital shadow ingests live data but does not complete the loop by issuing decisions or actuation based on the maintained state. A DT integrates both capabilities: it continuously assimilates measurements to maintain an actionable state and actively supports decision-making or control, subject to safety and security constraints.

A hierarchical DT framework for ITCS involves three tightly coupled workflows: a bidirectional data transmission that connects physical transportation systems to their digital representations, a modeling mechanism that specifies system states, parameters, and behaviors with high fidelity, and a synchronization loop that governs the calibration of DT states via measurements. **Data transmission workflow** spans vehicular onboard sensing, roadside sensing, and backhaul communication into edge devices with MEC and cloud services. Raw telemetry data are processed through feature extraction and alignment to produce time-aligned, quality-controlled inputs for further inference, prediction, and control. **Fidelity evaluation workflow** assesses vehicular geometry, dynamics, perception, communications, and human behavior representations and measures how well uncertainty is modeled and calibrated against observations. **Synchronization workflow** aligns the update frequency of DT with the decision horizon of downstream tasks. For autonomous vehicle (AV) control and collision avoidance, sub-second closed-loop operation with extremely low latency is required, whereas for traffic management, minute-scale updates are sufficient. In these cases, Age of Information (AoI) serves as a primary measure of staleness, where DTs must budget AoI relative to safety and performance constraints. Some common data assimilation mechanisms are used for synchronization, including state estimation for latent variables, parameter identification for slowly varying dynamics, and change detection for discrete events. After the initial construction, DT works as a continuous operating framework with integrated life-cycle management.

Taking ITCS as an example, end-to-end DT frameworks typically follow a mapping, optimization, and deployment pipeline. Mapping phase constructs machine-interpretable representations of roads, signals, demand, and network states; optimization phase derives policies that satisfy multiple objectives—such as safety, travel time, energy efficiency, and emissions—under operational constraints; and deployment phase integrates these policies into DTs for shadow-mode evaluation before releasing them into the field for controlled rollout. This pipeline leverages DTs both as a measurement substrate and as a safe-testing environment, thereby reducing risk and accelerating improvement. We adopt this pipeline structure to organize the discussion of different DT modules in ITCS from Sec. III to Sec. VIII.

B. DTs in ITCS

Autonomous and connected vehicles form a core pillar of DTs in ITCS. An SV-DT constructs a machine-interpretable state that includes geometry, kinematics, actuation limits, and health of the represented vehicle. It also maintains multiple executable models for perception, prediction, planning, and control that are compatible with live telemetry streams or recorded scenarios. Communication system extends the information horizon through V2X, aggregating updates from neighboring connected vehicles, roadside units, and infrastructure services. These connected components, which compose CV-DT and ENV-DT, also support predictive analysis, uncertainty quantification, and continuous monitoring.

Several key modules are required for SV-DT in practice. The perception module fuses multi-modal sensor measurements from cameras, LiDAR, and radar together from V2X components to infer individual objects, drivable space, road geometry, and traffic signal phase and timing. It processes these real-time telemetry streams and calibrates models to enable online adaptation to lighting, weather, and sensor aging, and to detect failures within the road environment. The prediction and planning modules connect perception data to physical transportation systems and policies. Short-horizon data-driven prediction estimates vehicle trajectories and interaction outcomes over seconds, while long-horizon time-series prediction captures intent, route choice, and cooperation patterns. The planning module converts these forecasts into a decision-making process, including feasible paths and control references subject to safety, security, and efficiency criteria. Here, DT also enables systematic what-if analysis by validating strategies under edge cases and supports counterfactual tests to explain decisions. Uncertainty is propagated through these stages to expose confidence and modulate caution in dense or adversarial traffic. Finally, the vehicle control module closes the loop, tracking controllers and stability modules translate planned references from DT into steering, throttle, and braking commands while respecting actuator dynamics. In DT, control policies are stress-tested across diverse scenarios before deployment, such as erratic human behaviors, variable road friction, and actuator faults.

Beyond the single vehicle, communications via V2X expand situational awareness and enable cooperative maneuvers with CV-DT. This type of DT emulates message exchange, congestion control, and handovers within large-scale vehicular network scenarios, continuously tracking end-to-end latency and communication reliability. With MEC integration on the roadside, perception or prediction tasks can be offloaded to edge servers near the vehicles to significantly shorten response times. Cooperative applications are emulated in CV-DT under realistic communication channels and load conditions. Second, integration with traffic management centers (TMC) further connects CV-DT to network-level objectives. DT streams aggregated metrics, such as queue length, arrival rate, and compliance levels, to adaptive signal controllers and variable speed advisory systems. In return, network-level control policies are formulated and validated in DT under measured demand and incident conditions prior to physical deployment.

On the physical infrastructure side, ITCS encompasses roads, intersections, signals, detectors, RSUs, and TMCs, particularly those equipped with intelligent controllers. These components maintain a consistent view of network topology, demand patterns, and incident states. The primary goals are monitoring and prediction at network scales, as well as serving as a reliable platform for evaluating operational policies before deployment. A growing body of datasets supports infrastructure-level ITCS-DT development. For instance, CitySim [15] provides drone-based vehicle trajectories with high-resolution mapping and signal timing.

C. Modeling and Learning Foundations for DTs in ITCS

ML supplies the predictive and decision-making capabilities that render a DT operational rather than merely descriptive. At a high level, ML models learn mappings from sensor telemetry streams transmitted over bandwidth-constrained wireless channels, and operational logs for state estimations, future forecasts, and action recommendations. Supervised learning utilizes labeled examples such as annotated images or ground-truth trajectories to train perception and prediction modules. Self-supervised learning reduces labeling effort by creating auxiliary tasks from raw data, such as reconstructing masked image regions or predicting next-motion segments. The output features usually transfer effectively across heterogeneous spatial conditions. Given that the transportation system is inherently multimodal, DT often fuses local sensor data (cameras, LiDAR, radar) with cooperative V2X messages and digital maps. Fusion architectures must explicitly account for communication constraints: late fusion combines outputs from specialized models to preserve bandwidth, while early fusion learns a common representation where one modality compensates for another during network degradation.

Temporal dynamics are critical in ITCS, as driving maneuvers, traffic flows, and wireless channel states naturally unfold as sequences. Consequently, models must capture dependencies across multiple horizons. Recurrent Neural Networks (RNNs) and temporal convolutions summarize short sequences for near-term forecasting of both vehicle trajectories and link qualities, while Transformer architectures handle long context windows and irregular sampling common in real-world network systems. Generative AI approaches, including diffusion and autoregressive models, synthesize realistic trajectories and scenes to rehearse rare or safety-critical situations scarce in historical data. Spatial structure is equally significant. Road networks and dynamic communication topologies are naturally represented as graphs, where Graph Neural Networks (GNNs) propagate information along lanes, intersections, and time-varying V2V/V2I communication links among vehicles.

Besides, translating predictions into actions requires policies that strictly adhere to safety and performance constraints, even under unreliable network conditions. Reinforcement learning (RL) optimizes policies through trial and feedback against explicit objectives such as delay, energy consumption, and communication overhead, while imitation learning distills expert behavior when exploration is costly. Model Predictive Control (MPC) [31] incorporates learned dynamics or cost functions to enforce feasibility at every step, explicitly compensating for V2X transmission latencies. Historical data from DT and the field enable offline RL, which improves policies without requiring additional on-road trials. Because operation on public roads demands quantified confidence, models estimate aleatoric and epistemic uncertainty via probabilistic outputs, ensembles, or Bayesian approximations. Calibration aligns predicted confidence with observed error to ensure meaningful thresholds, while conformal prediction offers coverage guarantees that translate into safety margins. Out-of-Distribution (OOD) detection and drift monitoring identify hidden failures when weather, sensors, traffic mixes, or network link qualities suddenly change. These signals inform the required DT synchronization frequency and determine the level of control authority granted to a policy.

Since conditions vary across cities, seasons, and hardware, ML models must adapt. Domain adaptation and sim-to-real transfer bridge the gap between rehearsal in DT and deployment on the road. Fine-tuning and meta-learning [32], [33] accelerate adaptation for new intersections or fleets, while continual learning updates models over time without catastrophic forgetting. When data cannot be centralized due to privacy or bandwidth limitations, Federated Learning (FL) usually trains shared models collaboratively across distributed agencies or fleets. This communication-efficient paradigm uses differential privacy and secure aggregation to protect vehicles' information while preserving utility over the network. Finally, decision support requires explainability and causal insight in ITCS. Attribution methods and counterfactual analysis reveal which inputs drove a perception or planning decision, while causal modeling separates correlation from effect to improve what-if analysis for policies such as variable speed limits. With these practices, ML components should respect DT synchronization budgets, minimize network payload, and transfer safely to closed-loop control.

III. Single-Vehicle DT (SV-DT)

SV-DT serves as a fundamental component within ITCS ecosystem. In this section, we survey the architecture of SV-DTs and examine their utility across a wide range of optimization algorithms and downstream tasks. While SV-DTs are often applied primarily in the context of Autonomous Driving System (ADS) control, our discussion incorporates broader perspectives—including driver behavior modeling and in-vehicle management—to provide a comprehensive view of vehicle-level control and optimization. Fig. 2 illustrates the core challenges and corresponding data-driven solutions encountered during the mapping, optimization, and deployment processes of SV-DTs. The extension of these concepts to connected systems, which enable network-wide coordination, will be subsequently discussed in Sec. IV.

A. Vehicle-State Mapping

1) Dynamic Virtual Representation: SV-DTs initially generate an accurate virtual replica by fusing historical and real-time onboard data, capturing kinematics, energy states (e.g., SoC, SoH), and driver behavior. Maximizing fidelity is critical for downstream tasks. This dynamic mirroring relies on time-series predictive analysis from vehicle logs [34]–[37],

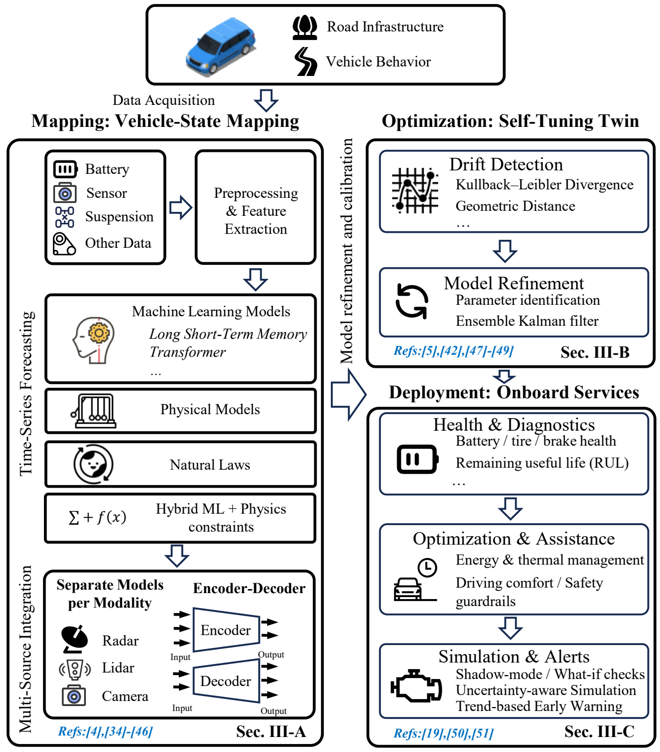

Fig. 2: Architecture of Single-Vehicle Digital Twin (SV-DT): from mapping, self-tuning, to applications.

extracting both short-term dynamics and long-term routines. While RNNs and convolutions are traditional, Transformers increasingly handle the extended contexts needed for continuous synchronization. These models act as evolutionary engines, keeping SV-DT current. For example, [38] creates a battery SV-DT pairing a fleet-trained SoH estimator with a continuously retrained SoC model. Similarly, [39] maintains driver-specific intent models for real-time edge predictions using LSTMs and inverse reinforcement learning. Such spatial-temporal pretraining and local adaptation sustain high-fidelity, living virtual representations.

2) Physics-Informed Constraints: Physics-informed constraints enhance SV-DT fidelity by ensuring state feasibility and reducing estimation bias. For vehicle tracking, [40] minimizes error efficiently by projecting poses onto planned paths and coupling physics-based lateral control with data-driven longitudinal regulation. Regarding thermal processes, [41] develops an SV-DT for a hydrogen fuel cell EV, modeling energy, HVAC, and battery dynamics across ambient conditions. Integrating physics improves ML interpretability, enforces natural and manufacturing limits, and facilitates realistic energy management evaluations.

3) Multi-Modal Data Fusion: Fusing heterogeneous, multi-modal onboard data (e.g., cameras, GPS, odometry) with varying sampling rates is essential. While late-fusion approaches process modalities separately [4], they incur high computational overhead. Conversely, encoder-decoder architectures learn shared feature spaces efficiently. For instance, a stacked Transformer fuses multi-channel contexts to capture diverse trajectories [42]. [43] uses a graph attention network and a conditional variational component on onboard perception to model agent interactions and generate future predictions. These architectures demand precise time alignment and uncertainty quantification. These uncertainty-aware representations form the time-stamped backbone for physical vehicle synchronization, enabling risk-aware rollouts and robust safety validation despite sensor imperfections [44]–[46].

B. Self-Tuning Twin Adaptation

1) Continuous Monitoring and Shift Detection: Following instantiation, SV-DT enters a continuous self-tuning phase that operates in parallel with onboard perception, planning, and control modules. The primary objective is to maintain alignment between SV-DTs and their physical counterparts without interrupting onboard service. Input data distributions inevitably drift due to evolving user habits, temporal variations, weather conditions, sensor aging, and component degradation. To address this, lightweight monitoring modules track distributions over salient features against recent baselines. Significant divergence or detected change-points trigger updates only when shifts are determined to be persistent rather than transient. To prevent model instability, updates are managed via strict governance protocols: confidence thresholds, cooldown timers, and canary evaluations that verify performance improvements prior to full adoption.

2) Parameter Identification and Data Assimilation: Upon confirmation of a distributional shift, two complementary mechanisms are employed to address distinct timescales in SV-DT. Parameter identification calibrates slowly varying coefficients within powertrain, thermal, and vehicle dynamics models, whereas data assimilation corrects fluctuating states in real time. Unified architectures integrate these processes, ensuring that latent states and physical parameters are refined simultaneously from shared observational data under a cohesive scheduling policy. For instance, [47] combines recursive identification with ensemble Kalman filtering to continuously refine SV-DT online. In operational settings, identification routines execute at lower frequencies with regularization to prevent overfitting, while assimilation tracks fast variables at the control loop cadence.

3) Reinforcement Learning for State Estimation: RL algorithms further augment self-tuning capabilities, particularly when vehicle operation states are difficult to measure directly or when the tuning process must adhere to complex constraints. In this case, RL policies are trained offline using historical operational data and subsequently deployed within SV-DT. For example, [5] demonstrates a battery management SV-DT where a deep deterministic policy gradient agent is trained on historical traces of voltage and current. Once deployed, this agent estimates SoC and SoH in real time. These estimates are then fed back into their SV-DTs to refine its internal monitoring state. This offline training, online deployment pattern adapts readily to other RL algorithms [42], [48], often incorporating runtime safety mechanisms such as action shielding and conservative updates to ensure that SV-DT's self-tuning remains within safe operational bounds.

4) Continual Learning for Adaptive Prediction: On the other hand, continual learning techniques complement identification, assimilation, and RL by updating learned perception and prediction components when only small batches of fresh data are available for SV-DT adaptation. The central challenge is preventing catastrophic forgetting by maintaining competence in established scenarios while adapting to new conditions. Methods fall into three families [49]. Rehearsal-based approaches maintain a small buffer of diverse past examples or generate synthetic samples for replay alongside new data. Regularization-based approaches penalize changes to parameters that were important for earlier tasks, preserving legacy behavior while allowing targeted adaptation. Architecture-based approaches allocate capacity dynamically using expandable layers, adapters, or masks, thereby isolating new knowledge or measurement from old.

C. SV-DT-enabled Onboard Services

1) Predictive Analysis and Health Monitoring: SV-DT transforms the mapped and self-tuned vehicle state (Sec. III-A) into an onboard service module that operates continuously alongside perception, planning, and control. The first service is predictive analysis and health monitoring, where real-time signals are fused with historical telemetry to forecast future SoC/SoH performance degradation, estimate remaining useful life, and anticipate upcoming failures so that maintenance can be scheduled proactively rather than reactively [19]. These forecasts are coupled with live key performance indicators and threshold policies, yielding early warnings for batteries, power electronics, tires, and brakes.

2) Embedded Digital Validation and Simulation: Beyond forecasting, SV-DT functions as an embedded digital validation platform for each vehicle, facilitating continuous simulation and scenario evaluation. High-fidelity vehicle and environment models are stress-tested under variable weather, road surface, load, and fault conditions to verify design assumptions and test operating envelopes without risk to the physical asset. Additionally, shadow-mode execution runs candidate control policies in parallel with the active production stack, logging divergences for analysis. Complementing this, counterfactual analysis investigates hypothetical outcomes by replaying historical scenarios with alternative control or calibration parameters. These capabilities enable rapid A/B comparisons and accelerate regression testing following updates. By significantly shortening the validation cycle before deploying SAFE-DT, SV-DT enhances system reliability and facilitates continuous improvement.

3) Decision-Support and Anomaly Detection: Optimization and decision-support services translate SV-DT state and predictions into actionable setpoints and policies [50], [51]. Energy management modules balance tractive power demands with thermal loads and accessory usage to extend range while preserving cabin comfort. Motion planning validation services score trajectory quality, stability, and passenger comfort, providing guardrail checks that gate handoffs between planning and control subsystems. Simultaneously, driver-centric services dynamically adapt assistance levels, alert timing, and human-machine interface layouts based on inferred intent and cognitive workload. These functions are designed to be compute-aware: lightweight solvers execute onboard at strict real-time rates, while computationally intensive analytics are deferred to maintenance or charging windows to respect thermal and power budgets. When the control system behavior deviates from nominal baselines, anomaly detection and diagnostic routines are automatically triggered. Residual tests, trend analysis, and learned detectors from SV-DTs identify deviations, localize likely root causes, and recommend corrective actions. Finally, the control and automation tier closes the loop with strict authority bounds, which will be further discussed in Sec. VII. This layered design enables the vehicle to benefit from adaptive strategies while preserving predictable behavior under uncertainty, providing auditable records that support after-action reviews and regulatory compliance [19].

4) Visualization and Data Governance: Finally, training, visualization, and collaboration services expose SV-DT to operators and engineers. Dashboards render twin states, confidence, and recommendations with clear rationales, and 3D replay tools support operator training, post-event analysis, and structured comparisons of calibration or policy changes as demonstrated in [50]. Governance enforces collection minimization, de-identification, and retention policies, while logging remains selective and event-triggered to conserve storage.

IV. Connected Vehicles DT (CV-DT)

CV-DT extends the virtualization paradigm from the vehicle unit to a collaborative vehicular network, as illustrated in Fig. 3. By integrating individual states from multiple collocated vehicles, a unified digital environment is constructed to facilitate cooperative perception, coordinated maneuvering, and collaborative decision-making. This framework scales from microscopic vehicle interactions to macroscopic traffic flow modeling within ITCS. Note that the following literature review focuses exclusively on vehicular dynamics and coordination strategies. The virtualization of the underlying infrastructure and communication is addressed separately in Sec. VI and Sec. VI. The ultimate goal of CV-DT is to provide a comprehensive, system-level analysis that addresses the limited view of SV-DT.

A. Cooperative Mapping & Multi-Source Data Fusion

1) Distributed Environmental Reconstruction: Building upon the foundational mapping and self-tuning capabilities of SV-DT, CV-DT elevates the abstraction level from a single-agent perspective to a shared, continuously updated environmental model. In this framework, observations from distributed vehicular nodes and infrastructure sensors are reconciled into a unified digital representation. This representation encompasses static elements (lanes, traffic controls, surface conditions) and dynamic context (drivable space, driver states) [52], [53]. Unlike the single-vehicle model, where mapping is intrinsically coupled to egocentric sensor suites and dynamics, CV-DT frames mapping as a distributed reconstruction process governed by explicit requirements for cross-participant consistency, common reference frames, and data freshness. This approach scales fundamental parameter identification and data assimilation by incorporating cross-source consensus mechanisms and large-scale parallel processing.

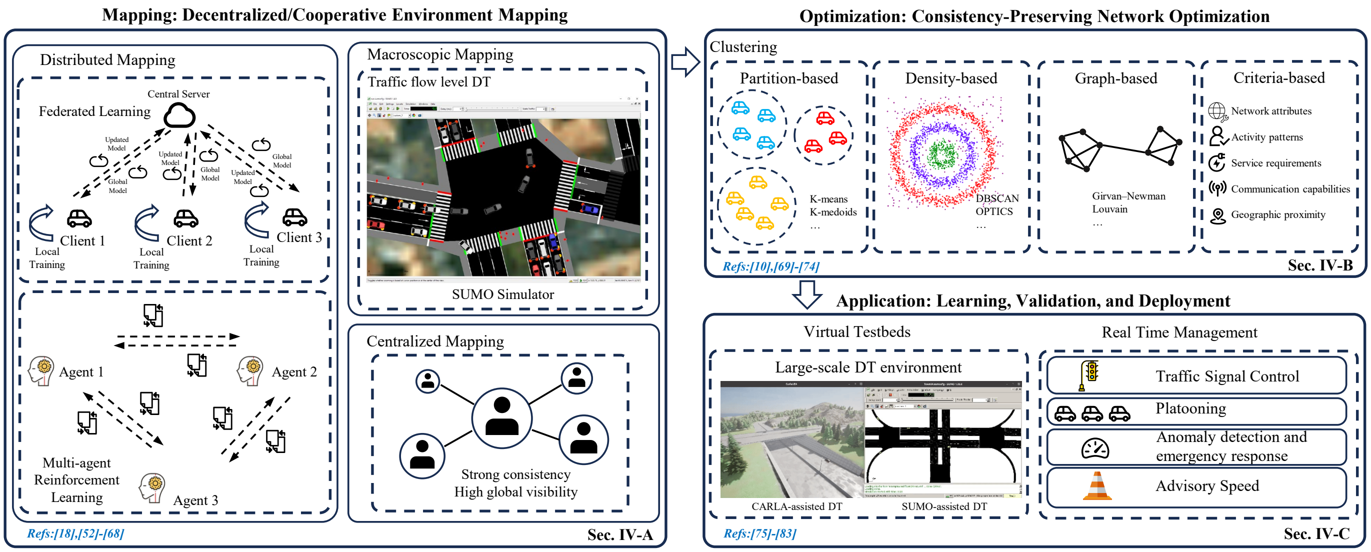

Fig. 3: Architecture of Connected Vehicle Digital Twin (CV-DT): from mapping, optimization, to applications.

2) Heterogeneous Data Alignment and Association: Constructing a coherent CV-DT from heterogeneous observations necessitates rigorous alignment and association. This process normalizes extracted features across diverse platforms, transfers calibration parameters between hardware variants, reconciles timestamps, and registers features into a common reference frame [54]–[56]. A crucial architectural principle is the separation of static content (e.g., lane centerlines, boundaries, curbs, signage) from dynamic artifacts (e.g., temporary blockages, work zones, debris). This segregation prevents transient objects from corrupting the base geometry. In multi-vehicle scenarios, overlapping coverage is consolidated to reduce map drift, utilizing confidence weighting and robust outlier rejection to ensure that a minority of corrupted reports cannot dominate updates [57]. While SV-DT relies on internal state correction to mitigate ego-drift, CV-DT extends this paradigm to cross-vehicle fusion, synthesizing multiple weak, partially overlapping perspectives into robust static hypotheses and transient dynamic overlays. Additionally, since mapping inputs vary in sampling rate, field of view, and reliability, efficient fusion requires compact, semantic descriptors rather than raw sensor data. Key primitives include vectorized lane segments with curvature attributes, landmark orientations, drivable-area polygons, and surface friction estimates derived from onboard slip cues [52], [53]. To balance responsiveness with stability, dynamic elements decay based on temporal validity, whereas static candidates are put into a permanent map only after receiving repeated, independent data streams [18]. Probabilistic occupancy grids [58] or evidential grids [59] often serve as intermediate representations for this extraction. Crucially, uncertainty is propagated alongside geometric features of connected vehicles, enabling downstream consumers to query not just the map geometry, but also its confidence levels and estimated staleness [60].

3) Decentralized Mapping: Beyond deterministic fusion, CV-DT increasingly relies on multi-agent, federated, or decentralized ML algorithms to maintain consistency without centralizing raw data. For instance, CV-DT typically adopts horizontal FL across similar participants to learn shared encoders, cluster-aware FL, and semi-synchronous aggregation to tolerate intermittent connectivity and heterogeneous compute while preserving timely updates [61], [62]. Personalization layers (e.g., lightweight adapters) [63], [64] specialize global backbones to local sensing stacks and operating regimes, while consensus- or gossip-based model averaging and edge aggregation at MEC nodes reduce bandwidth by merging updates near their source. Multi-agent learning, along with graph-aware analysis, ties local predictors to group objectives so that per-vehicle updates improve cluster-level consistency, with distributed risk-field and trajectory predictors serving as exemplars of collaborative inference under privacy constraints [65]. Secure aggregation and minimal sufficient statistics protect contributors, and drift detectors gate model pulls and pushes so stale or corrupted updates do not derail convergence. In combination, these decentralized methods align CV-DT mapping with realistic vehicular networks where data sovereignty, bandwidth limit, and asynchronous participation are the norm.

4) Network-Aware Ingestion and Edge Computing: Communication strategies and compute placement are intrinsic to the cooperative mapping pipeline. Consequently, CV-DT incorporates network-aware ingestion: contributions are aggregated and pre-validated at the network edge via MEC, filtering redundant data before it reaches the core. To optimize bandwidth, the system prioritizes updates that materially alter the map geometry, compressing payloads into semantic summaries and event-based deltas rather than raw streams [66], [67]. Instead of streaming dense raw imagery, participating vehicles publish only novelty keyframes, landmark updates, and concise topological ITCS. Metrics such as AoI and end-to-end latency are tracked on a per-tile basis, allowing the mapping service to dynamically throttle, defer, or accelerate ingestion based on current network contention and the urgency of the proposed ITCS [68]. To handle intermittent connectivity, opportunistic sidelink (V2V) exchanges disseminate local descriptors to maintain short-range consistency. Deferred uplinks then merge these cached ITCS once wide-area coverage is restored, ensuring map convergence without requiring simultaneous connectivity from all participants [54]–[56]. Further architectural details regarding network-wide DTs are provided in Sec. VI.

B. Consistency Preservation via Clustering

1) Efficiency and Fidelity through Clustering: Building on the cooperative mapping in CV-DT, consistency preservation increasingly relies on clustering to impose structure on large, heterogeneous fleets. Grouping vehicles that share mobility patterns, service requirements, sensing capabilities, or geographic proximity simultaneously enhances both the computational efficiency and the representational fidelity of CV-DT [69], [70]. Efficiency is improved by aggregating redundant data streams at the cluster level, which drastically reduces V2X bandwidth consumption and lowers the processing latency for state synchronization. Concurrently, fidelity is preserved because clustering restricts data fusion to statistically similar participants. This prevents low-quality or degraded sensors from corrupting high-precision cooperative maps, providing a natural aggregation unit for confidence weighting, outlier suppression, and localized validation before contributing ITCS to the global CV-DT state, as in [71].

2) Algorithmic Selection for Vehicular Clustering: Selecting a clustering method depends on data characteristics, real-time constraints, and semantics, especially in spatial-temporal vehicular networks [72], [73]. Partition-based approaches (e.g., k-means, k-medoids) are often favored for their computational efficiency and simplicity, particularly when cluster cardinalities are predetermined or can be estimated offline. Hierarchical methods, either agglomerative or divisive, are beneficial when capturing nested topological structures is essential, for instance, aggregating local neighborhoods into corridors, which subsequently form district-level clusters. Density-based algorithms (e.g., DBSCAN, OPTICS) excel at identifying clusters of arbitrary shapes and tolerating noise, effectively distinguishing sparse outliers from dense traffic streams. Graph-theoretic approaches, including spectral clustering and community detection (e.g., Girvan–Newman, Louvain), are ideal when relationships are naturally modeled as edges, such as shared routes, overlapping fields of view, or recurrent co-presence, where maximizing modularity reveals persistent cooperative patterns. Beyond algorithmic selection, the choice of extracted features is critical. Distinct partitioning schemes emerge from different attributes: mobility signatures (stop–go cadence, turn rates), sensor configurations (camera-only vs. multi-modal), or service affinities (common requests or constraints) all yield complementary structural insights [10].

3) Service-Oriented Hierarchical Structuring: Clustering serves not merely as a data-reduction mechanism, but as a foundational organizing principle for collaboration within CV-DT [69], [70]. In service-oriented architectures, vehicles requesting similar functionalities are organized into Service Request Groups. Each group is governed by an elected leader who negotiates with service providers on behalf of the collective members. This hierarchical structure streamlines communication, reduces negotiation complexity, and enables joint optimization over shared constraints and Quality of Service (QoS) targets [69], [70]. Leader election establishes a stable interface point for interaction with provider CV-DTs, while dynamic group membership facilitates the intra-cluster sharing of intermediate results (e.g., local map deltas, calibration parameters) and the fair allocation of limited resources, such as high-fidelity computational slots. This service-centric clustering integrates seamlessly with cooperative mapping by enabling group-level validation and prioritizing updates that maximize collective utility.

4) Grouped Training for Predictive Modeling: In CV-DT frameworks focused on autonomous operations, clustering provides the necessary structural substrate for distributed learning and optimization, directly enhancing DT's predictive capabilities. Instead of relying on a monolithic global model that struggles with concept drift, CV-DT utilizes grouped training and inference (e.g., cluster-based FL) to tailor predictive models—such as trajectory forecasting or intent recognition—to the specific behavioral dynamics of each cluster [71], [74]. By pooling statistically similar data distributions, grouped learning allows CV-DT to rapidly update its local virtual models to reflect immediate environmental changes while simultaneously reducing network contention by limiting broadcast scopes. When integrated with cooperative mapping, these clusters serve as the primary unit for staging local fusion, resolving low-confidence ITCS internally before escalation, and enforcing consistency policies that respect the distinct error characteristics and update cadences.

C. Learning, Validation, and Orchestration with CV-DT

1) High-Fidelity Virtual Testbed: Building upon the cooperative mapping and clustering, CV-DT now functions as a high-fidelity virtual testbed. It orchestrates complex, large-scale traffic scenarios by simulating diverse vehicle dynamics within a unified environmental context, enabling controlled, repeatable studies of interaction-dense behaviors [75]. By aggregating scene states, vehicle intentions, and environmental variables, CV-DT also provides a safe and scalable substrate for training data-driven models, particularly RL policies that benefit from dense interaction data [76]. This shared instrumentation further supports rigorous validation of V2X/V2V protocols, utilizing structured scenario libraries, counterfactual replays, and stress tests to systematically probe rare events and operational boundary conditions.

2) Policy Training and Strategic Optimization: Leveraging this virtualized environment, ML frameworks can be developed to optimize cooperative strategies for lane changing, trajectory planning, and system-wide traffic guidance. CV-DT serves as a primary generator for state observations and reward signals, strictly enforcing safety envelopes and operational constraints throughout the training process as studied in [75], [77]. Specifically, algorithms such as Deep Q-Learning and Soft Actor–Critic utilize extensive CV-DT rollouts to learn robust control policies capable of handling complex interactive dynamics. Complementing these functions, game-theoretic formulations introduce strategic awareness, modeling competitive and cooperative interactions to train policies toward Nash equilibria that remain stable even in dense traffic. Furthermore, inverse RL is employed to infer the latent objectives underlying human driving behavior, allowing CV-DT to align cooperative guidance with both individual driver preferences and broader network-level goals [77], [78]. Crucially, because CV-DT provides full state observability and deterministic reproducibility, offline policy evaluation and safety filters can be rigorously exercised at scale, validating performance before physical deployment.

3) Sim-to-Real Adaptation and Validation: To bridge the gap between simulation and real-world operation, CV-DT facilitates rigorous policy selection and calibration. This is achieved by replaying recent field scenarios against candidate controllers, quantifying performance gains using standardized metrics. Prior to active engagement, shadow-mode trials run candidate policies in parallel with the production stack, logging divergences without influencing live vehicle control [79], [80]. Policies distilled within CV-DT undergo sim-to-real adaptation via fine-tuning on local behavior distributions. Critically, their execution is guarded by rule-based shields—safety envelopes derived from conservative boundaries learned during simulation. When real-world operational drift is detected, CV-DT automatically updates its scenario libraries and triggers revalidation, effectively closing the loop between data curation, training, and field feedback. This cycle transforms CV-DT into a robust Continuous Integration and Validation pipeline for cooperative driving intelligence as studied in [81].

4) Real-Time Orchestration and Traffic Coordination: Transitioning from offline development to active deployment, CV-DT can also operate as an online orchestration layer. It supervises cooperative strategies in real time, coordinating interactions at critical conflict points and along corridors while compensating for sensing delays and actuation latencies. For instance, at non-signalized intersections, CV-DT facilitates cooperative passage via virtual slot reservation and consensus-based motion control, thereby minimizing unnecessary stops and resolving right-of-way conflicts deterministically [82]. To ensure scalability, hierarchical architectures partition responsibility: local CV-DTs manage neighborhood-scale coordination with low latency, while a global CV-DT reconciles objectives across broader regions. This layered approach ensures that system-wide efficiency and safety targets are maintained despite evolving traffic patterns [83].

V. Environment DT (ENV-DT)

ENV-DT constructs a high-fidelity digital replica of the physical operational domain, serving as the context for SV-DTs and CV-DTs. Its scope includes physical infrastructure assets such as roadway geometry, surface conditions, traffic control devices, and roadside sensor networks. Note that while communication infrastructure is essential for connectivity, its virtualization process is addressed as NET-DT, which is discussed in Sec. VI. By providing a comprehensive understanding of these external conditions, ENV-DT significantly enhances the situational awareness and safety of the broader vehicular ecosystem, as illustrated in Fig. 4.

A. Environment Reconstruction and Semantic Mapping

1) Photorealistic 3D Reconstruction: The foundation of ENV-DT lies in the geometrically accurate and realistic reconstruction of physical scenes. This process transforms heterogeneous sensor telemetry into DT-compatible physical assets that support downstream perception, planning, and policy evaluation. Multi-View Stereo pipelines [84]–[86] reconstruct dense point clouds or meshes from overlapping imagery, utilizing classical Structure-from-Motion (SfM) tools like COLMAP [87] or learned estimators such as MVSNet [88] and PatchMatchNet [89]. Complementing these, LiDAR-based pipelines contribute metrically precise range data, stabilizing scale and enhancing performance in texture-poor regions [90]. A paradigm shift in this ENV-DT reconstruction process is the emergence of Neural Radiance Fields (NeRF) [91]–[94]. NeRFs learn a continuous volumetric representation that maps spatial coordinates and viewing directions to color and density, enabling photorealistic view synthesis from arbitrary viewpoints. To fuse these modalities, LiDAR–camera fusion integrates the precise geometry of LiDAR with the enhanced semantic appearance of cameras [95], [96]. Other state-of-the-art architectures, such as BEVFusion [97], project multi-sensor evidence into a shared bird's-eye view representation, establishing a consistent geometric foundation for subsequent ENV-DT layers.

2) HD Mapping and Spatial Semantics: While geometric fidelity is essential, a raw 3D model remains insufficient for high-level ITCS decision-making in ENV-DT. An effective ENV-DT necessitates a structured, semantically rich layer that anchors assets to a stable coordinate system and exposes actionable, machine-readable context. Simultaneous localization and mapping [98] establishes the necessary spatial backbone by jointly estimating sensor trajectories and environmental maps. Visual-inertial and LiDAR-inertial frameworks, such as ORB-SLAM3 [99] and LIO-SAM [100], produce the consistent reference frames into which all subsequent semantic layers are anchored. Building upon this metric backbone, High-Definition (HD) mapping algorithms inject the detailed road semantics required for ENV-DT modeling and autonomous navigation. Deep learning models such as HDMapNet [101] and VectorMapNet [102] automatically vectorize sensor streams, extracting lane centerlines, road boundaries, crosswalks, and the attributes of traffic control devices. To capture complex relationships beyond mere geometry, graph-based environmental modeling organizes the scene into a queryable relational structure. In this formulation, nodes represent distinct objects or infrastructure elements, and edges encode spatial and semantic dependencies, thereby supporting reasoning tasks that demand deep contextual understanding rather than isolated detections.

3) Dynamic Behavioral Modeling: To transition from a static 3D reconstruction to a fully functional simulation environment, ENV-DT must be populated by dynamic behavioral models. This ensures that simulated assets evolve plausibly, subject to environmental rules and social constraints. Interaction-aware trajectory prediction models learn complex motion patterns from large-scale logs, producing short- and

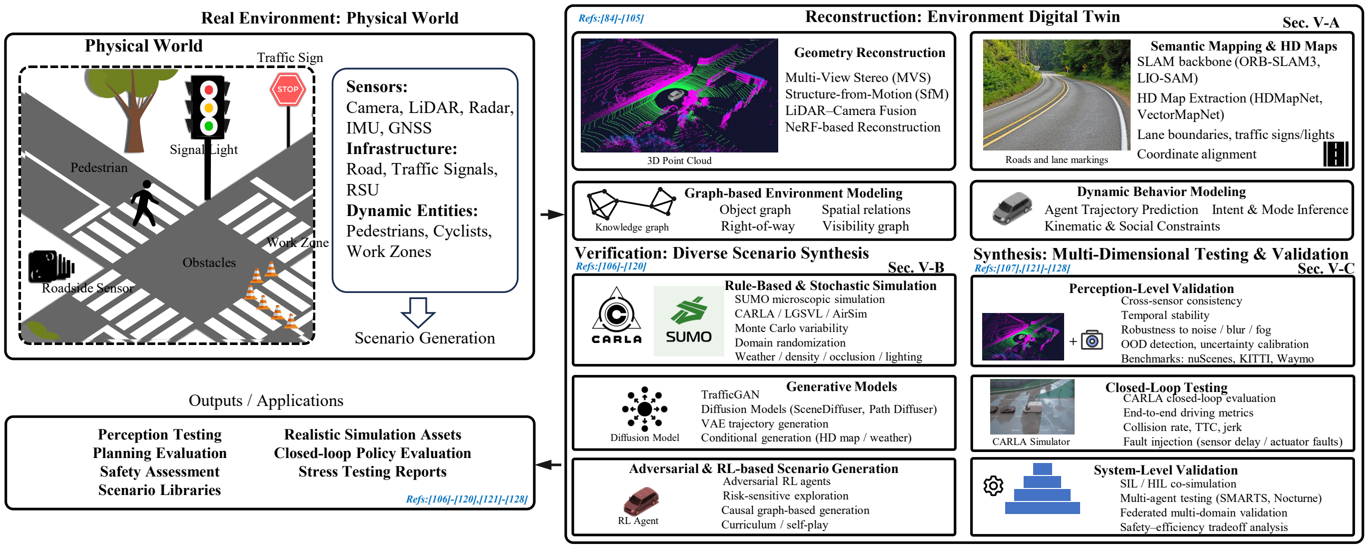

Fig. 4: Architecture of Environment Digital Twin (ENV-DT): from environment construction, scenario synthesis, to applications.

medium-horizon forecasts that account for scene context. Architectures such as Social GAN [103], Trajectron++ [104], and SceneTransformer [105] are pivotal here, as they explicitly model the latent dependencies between multiple agents. Ultimately, by integrating photorealistic reconstruction, structured semantic mapping, and predictive behavioral modeling, ENV-DT yields a comprehensive digital substrate, which is both visually faithful and operationally meaningful for validation and deployment of ITCS.

B. Diverse Scenario Synthesis

1) Stochastic and Rule-Based Generation: The conventional strategy for instantiating ENV-DT scenarios relies on rule-based and stochastic simulation and generation, where established engines generate road networks, vehicles, and sensor configurations before sampling variability to create diverse scene families. Traffic-level simulators like SUMO [106] excel at city-scale environment modeling, reproducing macroscopic flow phenomena and signal control policies. In contrast, high-fidelity platforms such as CARLA [107], LGSVL [108], and AirSim [109] focus on ego-centric realism, modeling precise vehicle dynamics and rendering physically accurate sensor streams like from camera, LiDAR, and radar. To ensure comprehensive coverage, diversity and uncertainty are introduced via two primary mechanisms. Monte Carlo sampling perturbs logical scenario parameters around a base configuration [110]. Simultaneously, domain randomization varies sensor and environmental appearance factors, including lighting conditions, textures, and noise models, to bridge the reality gap. The use of parameterized scenario templates ensures that this variation remains systematic and reproducible.

2) Data-Driven Generative Synthesis: While rule-based generation in ENV-DT offers explicit control, generative learning methods synthesize scenes by learning underlying distributions directly from real-world telemetry data. This ensures that the generated diversity from ENV-DT accurately reflects the structural complexity and statistical properties of physical environments. Generative Adversarial Networks (GANs) with scene-specific architectures, such as TrafficGAN [111], excel at producing plausible layouts and traffic snapshots. These models are typically conditional, generating content constrained by inputs like high-definition (HD) maps that define local topology and control assets. Diffusion Models have also emerged as the state-of-the-art for high-fidelity synthesis, operating by reversing a gradual noise-addition process. Variants such as SceneDiffuser [112] generate consistent static layouts, while Path Diffuser [113] synthesizes interacting vehicle trajectories that rigorously respect social norms and geometric constraints. On the other hand, Variational Autoencoders [114], [115] provide compact latent representations ideal for smooth, continuous multi-agent trajectory generation. Their latent space properties enable effective interpolation between driving regimes, facilitating targeted data augmentation near sparse or critical operating points. These generative frameworks afford high controllability by conditioning outputs on specific variables like map tiles, time of day, or weather. Evaluation has evolved beyond simple realism metrics to include task-grounded measures.

3) Adversarial Generation and Edge-Case Testing: To shift the testing focus from average operational cases to critical failure modes in ENV-DT, asset-driven and adversarial scenario generation techniques are employed. These methods train external agents to actively expose weaknesses in the system under test while strictly adhering to kinematic feasibility. Adversarial RL agents learn to manipulate secondary actors or environmental parameters to induce near-collisions, occlusions, or distributional shifts. Such strategies effectively concentrate training samples near the system's decision boundaries, revealing brittle interaction patterns. To ensure realism, frequent safety checks prevent these adversarial explorations from violating fundamental physical or traffic laws in ENV-DT [116]–[118]. Furthermore, Causal Graph-based composition [119], [120] imposes structural logic by encoding relationships among variables. This ensures generated scenes from ENV-DT obey cause–and–effect constraints, transforming ENV-DT from a mere content generator into a systematic safety validation instrument capable of certifying system robustness.

C. Multi-Dimensional Testing and Validations

1) Perception Consistency and Robustness: Validation within ENV-DT framework commences at the perception layer, aiming to certify that the reconstructed scene is geometrically consistent, temporally stable, and resilient to real-world perturbations. Cross-sensor consistency tests [121] verify intermodal alignment, ensuring that LiDAR point clouds, camera imagery, and radar returns map to coherent spatial elements after calibration. Concurrently, temporal consistency tests confirm that feature correspondences persist across consecutive frames despite viewpoint shifts or lighting variations. To assess resilience, robustness testing [122] actively probes failure modes by injecting controlled corruptions such as motion blur, sensor noise, adverse weather (fog, rain), and partial signal dropout, alongside simulated calibration drifts. Moving beyond raw accuracy metrics, the evaluation rigorously assesses uncertainty calibration, Out-of-Distribution (OOD) detection, and confidence-based gating mechanisms. This ensures that downstream planning modules receive predictions accompanied by quantifiable risk estimates. Finally, performance is benchmarked against large-scale, annotated datasets like nuScenes [123], Waymo Open [124], and KITTI [125]. These standardized baselines enable ablation studies to disentangle the impact of specific ENV-DT attributes—such as geometric fidelity, semantic richness, and sensor realism—on overall perception quality.

2) Closed-Loop Emulation and Stress Testing: Once perception is validated, the testing paradigm shifts to closed-loop emulation for decision-making and execution in ENV-DT. Unlike static log replay, ENV-DT ensures that the asset's actions actively influence the world state, subjecting planning and control stacks to realistic, dynamic consequences. High-fidelity platforms like CARLA [107], configured with verified ENV-DT assets (geometry, semantics, dynamics), facilitate rigorous end-to-end evaluation. Fault injection and stress testing further verify system resilience. By deliberately perturbing sensor timing, introducing actuator faults, or degrading surface friction, ENV-DT is probed for brittleness. These criteria focus on minimum risk maneuvers, bounded performance degradation, and autonomous recovery to nominal operation. Addressing the interactive nature of driving, multi-agent evaluation frameworks such as SMARTS [126] and Nocturne [127] generate dense, goal-driven traffic scenarios.

3) End-to-End System-Level Validation: System-level validation synthesizes perception, planning, control, and ENV-DT into comprehensive end-to-end experiments, designed to capture emergent behaviors within realistic operating settings. By combining Software-in-the-Loop (SIL) and Hardware-in-the-Loop (HIL) regimes, it exposes timing latencies and interface mismatches early in the design cycle. Furthermore, co-simulation couples vehicle dynamics with macroscopic traffic flow, ensuring that local maneuvers ripple through network-level conditions in a reproducible manner [128].

VI. Network DT (NET-DT)

NET-DT virtualizes the communication infrastructure and contexts that underpin the connected vehicle ecosystem, as illustrated in Fig. 5. By facilitating robust V2X information exchange, the communication system of ITCS enables cooperative perception, synchronized maneuvering, and collaborative decision-making that transcends the sensing limits of individual vehicles. However, reliance on connectivity introduces critical dependencies. The stochastic nature of wireless channels—manifesting as packet loss, jitter, and variable latency—can fundamentally compromise RSU-to-vehicle or V2V responsiveness. In safety-critical control loops, even minor transmission delays can destabilize vehicle operations and precipitate hazardous failures. Consequently, rigorous modeling of the communication network is indispensable. As a solution, NET-DT ensures that information sharing mechanisms are validated against realistic network impairments, guaranteeing the communication capability in ITCS and enhancing system-wide safety and performance.

A. Network-State Simulation

1) High-Fidelity Wireless Environment Representation: As shown in Fig. 5, network-state simulation within NET-DT framework constructs a time-aligned, high-fidelity virtual representation of the wireless environment [129], [130]. By evolving continuously with traffic and roadside infrastructure, it exposes critical cross-layer effects, tracing dependencies from physical channel properties and scheduling logic up to application-visible metrics like latency, jitter, and reliability. To achieve site-specific accuracy, NET-DTs employ a hybrid approach, blending measurement-driven learning with physics-based surrogates. This ensures that phenomena such as large-scale path loss, correlated shadowing, blockage, Doppler shifts, and co-channel interference are derived directly from the actual environmental geometry and mobility profiles. Deep learning architectures characterize these dynamics at multiple scales. For instance, CNN models analyze time–frequency channel response matrices to extract multipath structures for state classification [131]. Simultaneously, recurrent models (e.g., LSTMs, Transformers) forecast non-stationary load, link quality, and resource demand over horizons aligned with control and scheduling loops [10]. Given that vehicular connectivity forms a time-varying graph, GNNs are essential for propagating state information across dynamic topologies [132], [133]. By modeling vehicles and RSUs as nodes and links as edges, GNNs ensure that predictions of throughput and routing quality remain sensitive to neighborhood dynamics and mobility-induced churn. Collectively, these ingredients move beyond static analytic approximations, yielding NET-DTs that preserve causality and fidelity under real-world conditions.

2) Ray-Tracing and Hybrid Channel Modeling: Ray-tracing (RT) and hybrid physics–data models constitute the foundational pillars of NET-DT, particularly in regimes where propagation is heavily dictated by geometry and material properties. Deterministic RT methods, such as shooting-and-bouncing rays and image theory, compute path-level components anchored to high-fidelity 3D meshes of buildings, vehicles, foliage, and signage [134]. Image-based methods accelerate specular path computation, while particle or rasterization methods approximate diffuse scattering. These multipath components are synthesized into wideband/OFDM channel

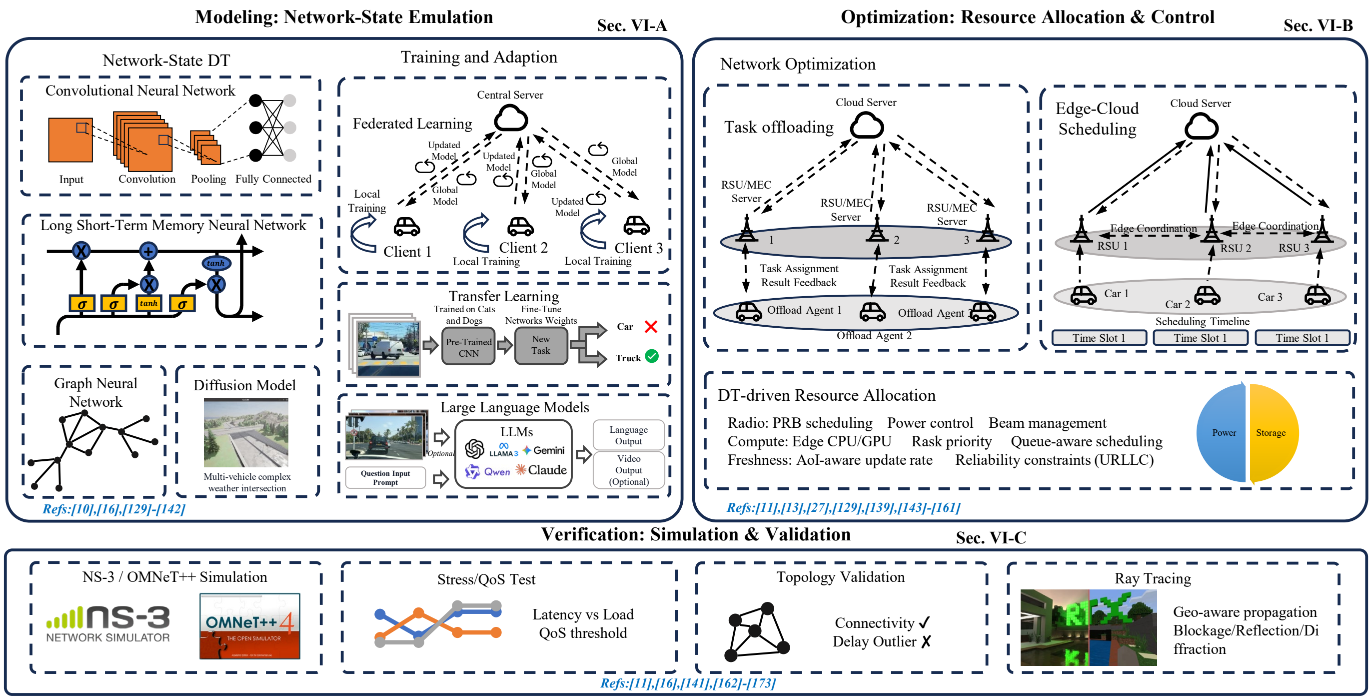

Fig. 5: Architecture of Network Digital Twin (NET-DT): from emulation, optimization, to verification.

impulse responses [135], which subsequently feed data link layers to drive network critical functions. To achieve city-scale coverage without prohibitive computation, multi-fidelity schemes spatially tessellate the domain: Computationally intensive RT is reserved for hotspots and complex intersections, while geometry-based stochastic channel models or 3GPP-style clustered delay line surrogates handle less critical regions [136]. Learned corrections bridge the gap between RT predictions for NET-DT and empirical measurements, while active learning algorithms intelligently select the next twin tiles or physical drive tests to maximally reduce NET-DT uncertainty. State-of-the-art frameworks integrate open-source RT backends, e.g., Sionna-based engines, with mobility models and NET-DT control loops as studied in [137]. This integration enables the study of spectral coexistence, per-resource-block interference, and beam dynamics across heterogeneous technologies like Dedicated Short-Range Communications and C-V2X/NR-V2X [16], [135].

3) Privacy-Preserving Twin Adaptation: FL is also pivotal for training and adapting NET-DTs without centralizing massive datasets, transcending simple local model training plus averaging paradigms to address complex network constraints [129], [138], [139]. In cross-device horizontal FL, numerous endpoints (vehicles, user terminals) collaboratively learn global encoders while retaining sensitive raw data (e.g., in-phase and quadrature samples, logs) locally. Techniques like client sampling, momentum correction, and model compression (sparsification/quantization) are essential to mitigate bandwidth and compute strain. Conversely, cross-silo horizontal FL [140] coordinates fewer, well-provisioned edge sites to train higher-capacity models on longer horizons. Vertical FL addresses scenarios where feature sets are partitioned across entities, preserving split learning that aligns entities without exposing raw feature columns, such as an infrastructure DT holding RF/contextual layers and a mobility model managing traffic semantics.

4) Decentralized Optimization and UAV Assistance: To enable timely decision-making under strict resource and privacy constraints, decentralized optimization serves as the runtime backbone complementing FL in NET-DT, transforming emulated network states into actionable control policies [130]. By decomposing complex joint control problems across vehicles, RSUs, and aerial relays, algorithms like alternating direction of multipliers allow each participant to solve a local subproblem [130]. The vehicles exchange low-rate messages to converge toward global feasibility and optimality, a process robust to intermittent links and heterogeneous compute capabilities. These two-stage frameworks first elicit truthful resource supply from infrastructure and demand from vehicles, then allocate tasks to maximize energy efficiency. In sparse coverage zones, unmanned aerial vehicle (UAV)-assisted dynamic NET-DTs mirror both network topology and task queues [141]. This enables rapid task reassignment and backpressure-based flow control, jointly stabilizing queues while optimizing energy consumption. Crucially, these optimizers operate in a closed loop with NET-DT, where constraints are parameterized by predictive analysis, while the real-world measurements feed back to recalibrate NET-DT.

5) Automated Orchestration and Validation via LLMs: Automation and rigorous validation complete NET-DT modeling block, significantly reducing the orchestration burden while ensuring that twins and control policies remain trustworthy [130]. LLMs can potentially function as intent-to-configuration translators [142]. For instance, given high-level goals (e.g., evaluate 30% V2X penetration during rain with a stalled vehicle event), they automatically synthesize coupled experiments involving ns-3 and SUMO mobility models. This includes selecting appropriate simulator variants, defining prediction horizons, generating ray-tracing jobs for new tiles, and producing data collection plans.

B. Efficient Calibration & Synchronization

1) MARL-Driven Continuous Calibration: At the control layer, NET-DT continuously calibrates its models and policies using global, time-aligned telemetry to surpass reactive heuristics and align with dynamic network conditions. Multi-agent RL (MARL) via the Centralized Training and Decentralized Execution (CTDE) paradigm is ideal here [11]. A NET-DT critic exploits system-wide observability during training, while lightweight vehicular and RSU actors execute policies locally. To scale effectively, value decomposition factorizes global rewards, mean-field/graph-based critics handle dense traffic, and safety shields constrain exploration to certified envelopes. These mechanisms navigate trade-offs like balancing NET-DT synchronization with primary tasks [143], or co-optimizing sensing, communication, and twin placement under strict constraints [144]. Complementing ML, game-theoretic designs manage self-interested agents: blockchain-secured contract markets reward high-quality data contributions [145], and Stackelberg formulations align caching incentives using trajectory-aware predictions [146]. Finally, recurrently co-located vehicles form collaborative caching clouds for efficient edge distribution [147].

2) Architectural Enablers and Network Slicing: From an architectural perspective, several enablers make such calibrations practical at scale for NET-DT. By strictly isolating ultra-reliable low-latency updates for NET-DT synchronization from high-bandwidth infotainment flows, it ensures that safety-critical freshness is maintained even under severe congestion. In this capacity, NET-DT executes admission control, elasticity scaling, and resource rebalancing by dynamically mapping per-slice key performance indicators to feasible radio, transport, and compute budgets [148], [149]. Further extending physical layer capabilities, Integrated Sensing and Communication (ISAC) shares waveforms and hardware between radar sensing and data transfer. Here, NET-DT scheduler dynamically optimizes the sensing–communication split and beam management. This minimizes AoI for environmental updates while simultaneously satisfying data rate and reliability targets [144]. Finally, adaptive offloading mechanisms and Deep Q-Network schedulers orchestrate the placement and chaining of these microservices in NET-DT. This minimizes end-to-end latency and energy consumption while preserving system resilience against node churn and infrastructure failures as validated in [150].

3) Transfer and Continual Learning at the Data Plane: At the data plane, efficient logistics for information and model transfer are vital. The architecture co-optimizes intra-twin synchronization and inter-twin exchange [151]. By predicting relay congestion, learner vehicles converge faster through knowledge transfer from expert agents, sharing few-shot information via optimal paths [152], [153]. Transfer and continual learning algorithms elevate NET-DT into reliable foundation models [129], enabling rapid adaptation to evolving corridors, spectrums, hardware, and policies. Standard methodologies pre-train on synthetic ray-traced corpora for universal radio priors, followed by fine-tuning on sparse in-situ measurements. To minimize edge compute overhead, parameter-efficient fine-tuning (e.g., adapters, low-rank updates) is employed [154]. Domain Adaptation aligns feature distributions to bridge environmental covariate shifts [155], while test-time adaptation tightens real-time calibration using self-supervised objectives without labeled data [27]. For long-term environmental evolution, continual learning frameworks prevent catastrophic forgetting using rehearsal buffers, weight regularization, and regional re-basing. Crucially, uncertainty quantification and drift detection gate model updates and throttle actions during low confidence [13], ensuring safe outputs. Collectively, these strategies deploy a single, resilient NET-DT stack that maintains accuracy across locales without full retraining.

4) Two-Tier Architecture and Hierarchical Scheduling: From the system perspective, to ensure calibration costs remain sustainable amidst evolving conditions, hardware, and policies, NET-DT lifecycle must support rapid adaptation. A two-tier learning architecture [156] addresses this by coupling cloud-based meta-learning with edge-side personalization. In the cloud, meta-models are trained across diverse, non-stationary spatiotemporal regimes to establish robust priors. At the edge, NET-DT executes fast, few-shot updates, capturing local idiosyncrasies, such as specific channel quirks, traffic rhythms, or hardware revisions, with minimal data requirements and downtime. Mirroring this structure, hierarchical RL optimizes resource scheduling within O-RAN-style virtualized environments [157]. High-level orchestrators determine strategic decisions like NET-DT placement, service modes, and synchronization strategies across network slices. Meanwhile, low-level controllers manage fine-grained primitives, utilizing asynchronous NET-DT updates to tolerate network variability and stragglers. Intelligent Reflecting Surfaces (IRS) actively reshape challenging propagation environments, while slice-aware pipelines maintain low-latency NET-DT refreshes in dense urban settings [149]. For high-dimensional optimization under strict deadlines, Quantum-aided MARL is employed for complex offloading and scheduling tasks [158], while a joint sensing–communication–placement policy is developed to optimize NET-DT locations to maximize both timeliness and energy efficiency [144].

5) Auditability, Safety Protocols, and Governance: To close the adaptive loop safely, calibration and evolution are governed by strict auditability protocols, ensuring updates improve performance without regressing certified behaviors. Transparency relies on exposing per-slice metrics and freshness. NET-DT enables counterfactual replays to attribute gains accurately, logging full dataset lineage, models, and configuration hashes for reproducibility. Deployment follows a rigorous path: policies are first evaluated in shadow modes, with subsequent live activation guarded by change-point monitors and automatic rollbacks. Smart contracts and provenance tracking underpin governance, ensuring fair data market compensation [145]. Concurrently, slice isolation and IRS-aided pipelines contain updates for safety-critical flows [149]. Ultimately, the synergy of learning-based orchestration [143], [144], slice/ISAC-aware architectures [148], [149], robust data migration logistics [139], [152], [153], [159]–[161], and rapid model adaptation [156]–[158] ensures NET-DT operates as a fresh, efficient, and verifiably safe closed-loop system under dynamics.

C. QoS/QoE Assurance Applications

1) Predictive QoS Assurance and Telemetry Fusion: NET-DT framework enables proactive assurance of Quality of Service (QoS) and Quality of Experience (QoE) by combining real-time monitoring with short- and medium-horizon prediction, allowing network resources to be dynamically steered before congestion, signal blockage, or outages materialize in ITCS. Time-aligned telemetry from vehicles and infrastructure is fused to continuously track critical metrics, including trajectories, queue backlogs, channel states, and beam dynamics. Extended Kalman filters support low-latency motion tracking for precise assignment and beam planning, while sequence models forecast mobility patterns and network load to anticipate hotspots, handover pressure, and interference growth [16], [162]. Augmenting these data-driven predictions, NET-DT-assisted ray-traced radio maps provide site-specific forecasts of link quality and interference that account for changing geometry and traffic, enabling advanced strategies like preemptive spectrum reuse and precise power control [16].

2) Dynamic Routing and Resource Offloading: Building on its predictive foundation, NET-DT acts as a simulation testbed and runtime supervisor for complex routing, offloading, caching, and scheduling. For multi-hop routing, frameworks like ELITE [141] fuse parallel single-objective learners to dynamically select resilient routes under non-stationary channel conditions. For task offloading across the vehicle-RSU-cloud continuum, cluster-aware multi-agent learning reduces search space complexity by aggregating vehicles locally before decentralized scheduling [163]. Toward unified management, VECO [164] co-optimizes predictive caching with RL-based offloading to logically co-locate storage and compute demands. Enhancing edge intelligence, NET-DTs utilize knowledge graphs and Double Deep Q-Networks to compress state representations, accelerating multi-hop decision-making [165]. Finally, dual-layer architectures minimize non-convex end-to-end collaborative task delays via alternating optimization of compute frequency, bandwidth, power, and offload ratios [166].

To ensure consistent QoS under high mobility and signal blockage, NET-DT orchestrates joint communication–sensing and spectrum intelligence. Using ISAC, NET-DT leverages motion prediction for vehicle-to-RSU assignment and predictive beamforming, maximizing communication rates and sensing accuracy [162]. Extending to visible light domains, NET-DT minimizes latency by decomposing communication and sensing sub-tasks via alternating minimization [167]. At the spectrum layer, the PRISM framework [16] couples ray-traced channel predictions with position tracking to cluster V2I links and allocate transmit power via Multi-Agent Q-Learning. Complementing this, RAVEN [168] combines channel prediction with graph-theoretic clustering to optimize multi-hop millimeter-wave capacity and mitigate co-channel interference. Finally, IRS and slice-aware update pipelines isolate and stabilize NET-DT synchronization flows in dense urban environments [169].

3) Intent-Driven QoE Optimization and Task Scheduling: QoE assurance optimizes user-perceived utility over raw link metrics. NET-DT achieves this by mapping technical indicators into a unified satisfaction score to guide intent-driven resource allocation across network slices [170]. For scalability, coordination graphs restrict multi-agent dependencies, enabling locally consistent decisions. Graph-assisted NET-DTs with Lyapunov objectives effectively stabilize queues and energy while meeting strict task deadlines [171]. In dynamic terrestrial–UAV environments, combining Asynchronous Federated Learning with MARL optimizes robust task offloading by balancing completion rates, energy, and delay [11]. For complex tasks, parallel intelligence frameworks schedule directed acyclic graphs across the vehicle-edge-cloud continuum, utilizing auction-based allocation and evolutionary optimizers to satisfy precedence and latency constraints [172]. Finally, closed-loop NET-DT feedback dynamically adapts online control parameters, such as re-tuning partial offloading in response to real-time mobility and load shifts [173].

VII. Safety & Security DT (SAFE-DT)

SAFE-DT is illustrated in Fig. 6, which extends the virtualization paradigm to rigorously encapsulate the risk profiles and protective mechanisms of ITCS. This section discusses safety and security from two perspectives: first, securing the inherent architecture of DT itself against compromise; and second, utilizing SAFE-DT as a sandbox to model and assess the safety and security of physical ITCS. The primary objective, however, is to leverage SAFE-DT to explicitly represent potential vulnerabilities, catastrophic failure modes, and dynamic threat responses within the physical domain. By providing a high-fidelity environment for proactive risk assessment, SAFE-DT becomes indispensable for ensuring ITCS resilience and certifying the reliability of safety-critical applications.

A. Trustworthy Twinning Process

1) Cryptographic Identity and Authentication: A trustworthy twinning process requires cryptographically anchoring physical vehicles to their DTs to prevent identity spoofing and Sybil attacks that inject false data or disrupt control loops. At the intra-twin layer, multi-factor authentication ensures definitive ownership and liveness prior to sensitive exchanges [174], [175]. Implementations leverage efficient elliptic-curve cryptography (ECC), alongside proxy signatures for delegated fleet operations and Key-Policy Attribute-Based Encryption, to strictly align decryption capabilities with specific roles and contexts [174], [175]. To thwart replay, relay, and downgrade attacks, session bootstrapping enforces nonce challenges, certificate pinning, and freshness tokens, utilizing short-lived tickets to minimize compromised key exposure. Beyond software, hardware-backed roots of trust provide immutable identity. This enables SAFE-DT to use remote attestation to verify unmodified in-vehicle firmware before granting privileges, mitigating rootkit risks. Protocol correctness is rigorously validated via BAN logic, AVISPA model checking, and random-oracle arguments, providing formal assurance for safety-critical deployments [174], [175].

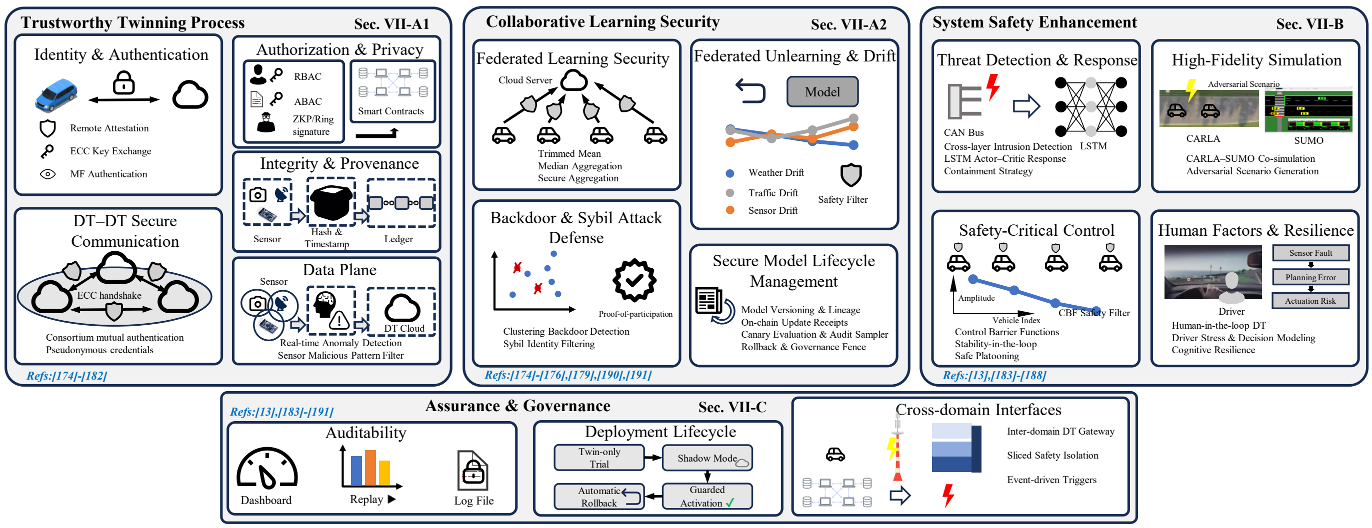

Fig. 6: Architecture of the Safety and Security Digital Twin (SAFE-DT): from twinning process, safety enhancement, to performace guarantee & governance.

Finally, inter-DT exchanges spanning organizational boundaries employ consortium-style ECC mutual authentication. This balances anonymity with operational accountability by utilizing pseudonymous credentials [176].

2) Blockchain-Enabled Provenance and Integrity: To combat data falsification during the twinning process, integrity and provenance mechanisms establish a tamper-evident chain of custody throughout SAFE-DT lifecycle. Consortium blockchains act as an immutable ledger, recording state transitions, sensor digests, and policy decisions. Hashing these events ensures non-repudiation and enables decentralized cross-party audits [176], [177]. To prevent unauthorized tracking, pseudonym systems decouple long-term credentials from operational addresses. Synchronous group pseudonym changes, orchestrated via blockchain, eliminate timing artifacts to thwart linkage attacks as vehicles cross RSU boundaries [178]. To support real-time SAFE-DT control, ledger scalability is optimized via transaction batching and state channels for high-rate telemetry. Layered rollups post summarized proofs to the main chain, balancing full auditability with the low latency required for safety loops. In self-organizing overlays, signed messaging, trust-aware forwarding, and message fragmentation limit the impact of compromised relays, while alias exchanges and scope-limited credentials minimize metadata leakage [177]. Finally, safety-grade key management (e.g., hardware-backed generation, rapid rotation) and continuous event mirroring in DT ledger maintain a forensic chain of custody for all system safety configurations.

3) Data Plane Defense and Sensor Reconciliation: Establishing trust within the data plane demands a defense-in-depth strategy that blends cryptographic immutability with behavior-aware scrutiny and cross-source reconciliation. To thwart retroactive tampering, sensor streams are hashed and timestamped at the point of collection, with digests anchored to the ledger. Channel-local buffering retains signed preimages, enabling on-demand challenge–response verification. To detect false data injection and spoofing attacks, DT-resident anomaly detectors learn dynamics-consistent manifolds. These ensembles utilize calibrated uncertainty and physical limit checks to reliably distinguish benign noise or driver error from malicious inputs as proposed in [179]. Cross-sensor reconciliation further hardens ITCS against discordant reports by aligning inertial, odometric, and visual cues. Map priors and signal phase constraints act as sanity checks, automatically downweighting physically impossible trajectories. For full auditability, provenance chains track every data transformation—calibration, filtering, fusion, inference. This allows auditors to reconstruct exactly which inputs and model versions produced a specific decision. To preserve availability without compromising safety, actions are downgraded to advisory status when confidence falls below thresholds. SAFE-DT then requests corroboration from redundant channels, preventing untrusted evidence from driving physical actuation. Finally, for investigative purposes, privacy-preserving watermarks enable the tracing of leaked datasets without exposing sensitive identifiers in operational streams [180], [181].

Deploying fleet-scale collaborative intelligence in ITCS inevitably expands the attack surface. Within SAFE-DT framework, federated learning pipelines are rigorously hardened using Secure Aggregation protocols, ensuring servers observe only encrypted updates to preserve privacy. Concurrently, Robust Aggregation dampens malicious inputs and dynamically isolates suspicious clients. SAFE-DT also manages client attestation to bind dataset integrity to hardware trust roots, while Differential Privacy bounds inversion risk without sacrificing control utility [176], [179]. To counter distributed threats like model replacement, backdoors, and Sybil attacks, SAFE-DT orchestrates update norm clipping, gradient anomaly scoring, and proof-of-participation tokens. Addressing non-IID vehicular data, SAFE-DT utilizes geography- or channel-based clustered aggregation to reduce gradient conflicts, alongside personalization layers that adapt backbones to local regimes. For accountability, on-chain update receipts bind contributions to revocable pseudonyms, enabling traceable rollbacks without exposing driver identities [176]. Audit samplers periodically replay updates against canary datasets to detect silent drift, automatically fencing suspect contributors. By integrating these measures with foundational authentication, authorization, integrity, and provenance [174], [175], [178], [182], SAFE-DT elevates the twinning process into a verifiable, resilient cyber–physical workflow, ensuring ITCS intelligence remains auditable and aligned with systemic safety and privacy obligations.

B. System-Level Safety Enhancement

1) Threat Detection and Fail-Operational Control: SAFE-DT framework enhances ITCS safety by transforming threat anticipation, detection, and mitigation into a continuous, closed-loop process that explicitly reasons about uncertainty, complexity, and drivers' behavior across both cyber and physical domains. In the cybersecurity domain, DT-resident orchestrators utilize Long Short-Term Memory (LSTM) Actor–Critic agents to identify spatiotemporal attack signatures that elude static rules. These agents capture subtle cross-subsystem correlations, such as powertrain anomalies synchronized with suspicious bus traffic, triggering layered containment strategies with calibrated false-alarm rates [183]. To preserve trustworthy logging and coordination under conditions of mobility and partial connectivity, blockchain consensus is virtualized within SAFE-DT's cyberspace. Mechanisms like Diffused delegated Byzantine fault tolerance reduce confirmation latency for safety-relevant events, maintaining liveness during handovers while dynamically adjusting membership stakes to reflect the real-time availability of edge servers [184]. On-road safety is rigorously enforced through robust traffic control protocols. Control barrier function constraints are embedded directly into DT-in-the-loop controllers, preserving lane-level invariants and string stability despite communication delays. Calibration processes utilize goodness-of-fit and measure-of-performance metrics to strictly align plant models with field behavior prior to release [13]. Furthermore, dynamical complexity in mixed traffic is managed explicitly: by modeling heterogeneous human and automation responses, SAFE-DT identifies control gain regions that avoid bifurcation and chaos. This bounds butterfly effects that could otherwise amplify small perturbations into catastrophic prediction failures as studied in [185]. Finally, runtime assurance architectures pair high-performance learners with certified safety governors. These governors intervene only when predictions stray outside verified envelopes, ensuring fail-operational behavior rather than simple fail-safe shutdowns.

2) Human-Centric Monitoring and Systemic Resilience: In this context, human factors and systemic resilience are treated as safety-critical components rather than exogenous disturbances. SAFE-DT monitors driver state in real time by classifying dynamics such as steering variance, acceleration signatures, and lane-keeping fidelity into risk levels and routes targeted interventions via augmented reality heads-up displays that place salient cues within the driver's field of view without inducing distraction [186]. Traveler response models quantify rerouting, compliance, and demand elasticity under disruptions; SAFE-DT uses these to forecast secondary effects such as spillover congestion and to design mitigations that do not merely displace risk in space or time as studied in [187]. At subsystem level, correlation and causality analysis map how sensor faults or cyber intrusions propagate through perception, planning, and actuation, distinguishing random failures from coordinated attacks and selecting the least-regret fallback strategy, e.g., degraded autonomy, advisory-only outputs, or controlled stop, while preserving essential services [188]. Crew-in-the-loop drills and tabletop exercises run inside SAFE-DT so that operators rehearse incident playbooks on realistic reconstructions of past events, which increases readiness and shortens response times when similar patterns reappear.

C. Performance Guarantee and Governance

1) Virtual Validation and ODD Coverage: To validate safety at scale without exposing public roads to risk, high-fidelity scene understanding supplies calibrated, reproducible environments. By stitching aerial LiDAR, map assets, and ego-vehicle sensors, ITCS constructs lane-accurate geometry replete with conflict zones and priority rules. This virtual substrate integrates CARLA for micro-dynamics and sensor rendering, SUMO for macroscopic traffic flow, and physics engines to ensure contact realism for collision envelopes [189]. Within this testbed, SAFE-DT curates comprehensive scenario libraries spanning nominal, rare, and adversarial cases—such as occlusions, sudden cut-ins, degraded friction, or emergency priority events. It executes accelerated what-if campaigns to rigorously evaluate vehicle–object interactions under controlled perturbations [189]. Crucially, validation surpasses open-loop accuracy checks by enforcing closed-loop outcome metrics, including intervention frequency and safety margin distributions. Acceptance criteria fuse classical risk indicators (e.g., Time-to-Collision, Time Headway) with DT-specific measures, such as data freshness distributions, decision latency budgets, and actuator saturation margins, ensuring controllers strictly respect certified envelopes prior to field activation [13]. Finally, systematic scenario sampling from SAFE-DT ensures proportional representation across geographic, weather, and traffic regimes to satisfy operational metrics.

2) Adaptive Resilience and Federated Unlearning: Based on this pipeline, adaptive resilience extends deep into a learning substrate, ensuring predictive engines within SAFE-DT remain secure under concept drift, distribution shift, and adversarial pressure. When federated training is compromised by polluted participants, federated unlearning offers a surgical remedy [190]. Orchestrated by SAFE-DT, these unlearning agents identify and excise malicious influence without prohibitive full retraining, utilizing targeted gradient corrections and dependency tracking to roll back tainted contributions while preserving legitimate learning. Concurrently, detection pipelines evolve via Stacked Sparse Autoencoders, which compress high-dimensional telemetry into informative manifolds. An online learner then dynamically selects the optimal classifier for the current regime, significantly improving recall on emerging attack families within edge latency budgets [191]. Distribution shift monitors continuously analyze weather patterns, fleet mix, and sensor health. When guard bands are exceeded, these monitors trigger spaced recalibration or policy retuning. Furthermore, control-side safety shields enforce rigorous action filters derived from counterexamples generated during adversarial scenario search. This ensures that even partially degraded perception pipelines cannot yield unsafe actuation sequences, effectively bounding risk during periods of adaptation.

3) Lifecycle Governance and Automated Rollback: As the last step, governance frameworks close the safety loop by rigorously linking DT-derived decisions to auditable artifacts and enforcing staged activation protocols with automated rollback capabilities. SAFE-DT platform in [13], [184] exposes real-time safety and freshness key performance indicators per network slice. It facilitates counterfactual replays to attribute performance gains to specific mitigations, such as routing adjustments, barrier-function tightening, or consensus tuning, while maintaining immutable logs of inputs, model versions, and calibration snapshots to enable rigorous after-action review and rapid rollback. Policy deployment follows a graduated path: initiating with twin-only simulation trials, advancing to shadow modes that log divergences without influencing active control, and culminating in guarded activation. During this phase, change-point monitors stand ready to trigger automatic rollback upon detecting distribution shifts or anomaly bursts. Crucially, cross-domain interfaces, including vehicle-to-grid power quality protections, slice isolation for safety messaging, and incident-triggered consensus escalation, are validated against comprehensive scenario catalogs. This ensures that optimizations in one subsystem do not inadvertently degrade performance in another [183], [184]. Safety cases are maintained as living documents, dynamically linking safety claims to empirical SAFE-DT evidence, e.g., traceable metrics, scenario coverage, and replayable experiments, thereby streamlining external audits and regulatory reviews without the need for bespoke physical test rigs. Ultimately, the integration of high-fidelity scene generation [189], provable safety control [13], dynamical complexity analysis [185], behavior-aware monitoring [186]–[188], and self-healing learning pipelines [190], [191] transforms SAFE-DT into a continuous, verifiable safety engine, equipping ITCS with robust resilience against cyber threats, human variability, and deep environmental uncertainty.

VIII. Twin of Twins (ToTs): An Integral Framework

Leveraging the high-fidelity, domain-specific DTs detailed in the preceding sections, the framework culminates in the integration of these discrete components into a comprehensive, macro-scale DT architecture. This paradigm is termed Twin of Twins (ToT)—a city-scale System of Systems constructed by synthesizing the virtual substrates of individual vehicles (SV-DT), cooperative vehicular fleets (CV-DT), physical infrastructures and environments (ENV-DT), communication networks (NET-DT), and overarching safety and security constraints (SAFE-DT). This unified integration thereby enables the holistic optimization of urban mobility and critical downstream sectors such as energy grids and emergency response. Achieving this synthesis requires the rigorous orchestration of both inter- and intra-DT operations, ensuring seamless semantic alignment, data exchange, and functional interoperability across heterogeneous domains. Crucially, the architecture must sustain real-time responsiveness across this massive integrated framework, supporting dynamic decision-making while preserving overall system reliability and effectiveness at scale.

A. Heterogeneous Twin Integration

1) Modular Co-Simulation and Variable Granularity: The evolution of DT technology is defined by the shift from isolated, single-purpose models to a holistic orchestration of the entire transportation ecosystem. A practical implementation of this vision relies on *modular* co-simulation, employing standards like Functional Mockup Interface to synthesize composite, high-fidelity models by linking best-in-class, domain-specific simulators—ranging from vehicle dynamics to network communication stacks [192]. This architectural principle underpins ToTs paradigm, which aggregates multiple task-oriented DTs into a unified, interoperable system model. Prominent implementations include architectures that couple high-fidelity vehicle dynamics (e.g., IPG CarMaker) with microscopic traffic flows for comprehensive safety analysis [12], and open-source frameworks like VaN3Twin that fuse mobility simulation with advanced ray-tracing to model V2X channels with physical precision [135]. However, the inherent complexity of such large-scale integrations demands rigorous efficiency management. Consequently, variable-granularity ToT frameworks have emerged, utilizing multi-objective RL to dynamically adjust the fidelity and update frequency of constituent twins. This strategy optimizes the trade-off between system utility and resource consumption as in [193].

2) Bidirectional Data Exchange and Vehicle-in-the-Loop: Specifically, the efficacy of any integrated twin framework relies on a robust data backbone that facilitates tight, closed-loop coupling with the physical system. A defining characteristic of these advanced systems is bidirectional data exchange: telemetry flows from physical assets to the virtual domain for analysis, while actionable intelligence flows back to physical actuators to guide operations. Early iterations demonstrated this utility in targeted applications like cooperative ramp merging, where real vehicles uploaded state data to cloud-based DTs to receive computed advisory speeds [194], [195]. Modern implementations have evolved to construct comprehensive, real-time mirrors of entire traffic corridors by synchronizing diverse V2X data streams, including Signal Phase and Timing (SPaT), MAP data, and Basic Safety Messages (BSM) [196]. This continuous fusion ensures that DTs strictly replicate physical traffic states, enabling the feedback of validated advisories, such as speed recommendations or incident notifications, directly into live traffic management systems [196]. Furthermore, this pipeline underpins Vehicle-in-the-Loop (ViL) testing, allowing physical test vehicles to interact with virtual objects via V2X. This capability is indispensable for safely validating hazardous scenarios that are impractical to reproduce on public roads [14].

3) Human-in-the-Loop and Mixed Reality Integration: On the other hand, a comprehensive ToTs integration must inevitably account for the most stochastic element of ITCS: the human actor. Modern frameworks are increasingly designed to support pedestrian- and driver-in-the-loop simulations to create a more realistic and effective testing environment. To

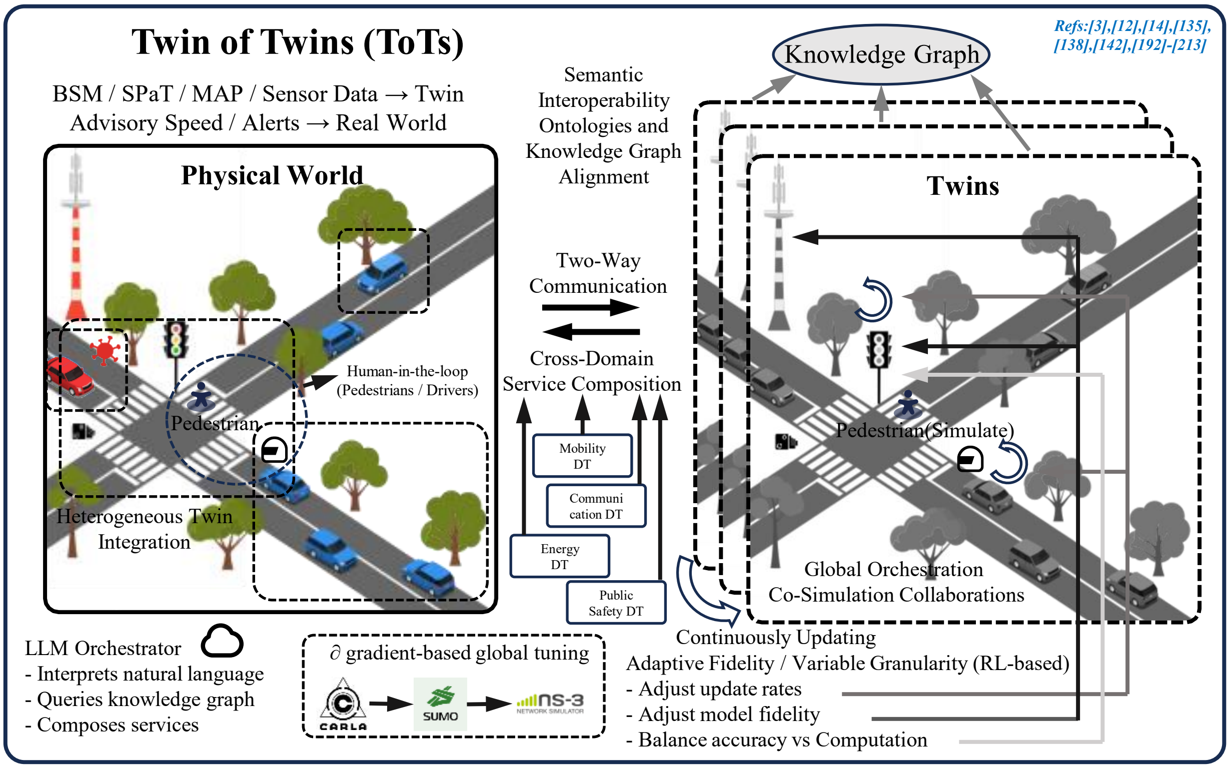

Fig. 7: Architecture of Twin of Twins (ToTs): from twin integration, global orchestration, to cross-domain services.

integrate pedestrians, ToTs architecture employs a dual approach. Passive integration utilizes edge computing to process surveillance feeds into lightweight, semantic agent models, preserving privacy while capturing naturalistic behavior [3]. Conversely, active integration immerses real subjects within a Cave Automatic Virtual Environment (CAVE), enabling real-time interaction with simulated traffic, i.e., a critical capability for validating Vehicle-to-Pedestrian (V2P) collision warning systems [197]. For driver integration, the Mixed DT concept establishes a unified testbed where physical miniature vehicles, purely virtual entities, and human-driven actors—augmented via Mixed Reality interfaces like HoloLens—coexist within a synchronized state space [198]. Orchestrating these multi-modal systems demands advanced AI techniques; multi-task and multi-modal learning algorithms process diverse data streams simultaneously to maintain coherence. Furthermore, transfer learning and fine-tuning are essential to bridge the reality gap, adapting models trained in sterile simulations to the noisy, complex environments characterized by genuine human participation [199]–[201].

B. Global Orchestration & Co-Simulation

1) Multi-Agent Twin Orchestration: Orchestrating a city-scale DT ecosystem demands robust distributed intelligence paradigms capable of managing vast, heterogeneous agent populations and fragmented data silos. Federated multi-task learning emerges as the foundational methodology, facilitating collaborative model training across distinct entities while strictly preserving data sovereignty. This collaboration manifests in two forms: Horizontal FL, where twins (e.g., from different municipalities) sharing identical feature spaces collaboratively refine a unified model; and Vertical FL, where twins possessing disparate feature spaces but overlapping user entities (e.g., dynamic vehicle telemetry vs. static infrastructure logs) jointly synthesize comprehensive predictive models. To mitigate the specific challenges of intermittent connectivity and heterogeneous computational power in vehicular networks, asynchronous FL frameworks are essential. Unlike synchronous protocols that are bottlenecked by the slowest participants (stragglers), asynchronous approaches permit dynamic global model updates, thereby significantly enhancing training efficiency and scalability as studied in [138]. Complementing this collaborative learning is Multi-Agent Reinforcement Learning (MARL), which governs collaborative decision-making. Typically implemented via the Centralized Training with Decentralized Execution (CTDE) paradigm, this approach utilizes DT's global visibility to train sophisticated agent policies offline, while enabling individual agents to execute tasks online based solely on local observations. This architecture proves highly effective for complex joint optimization, such as concurrently tuning vehicular sensing strategies and DT resource deployment to maximize system-wide QoS [202].

2) Differentiable Co-Simulation and Semantic Interoperability: For implementation, a critical technical challenge in realizing ToTs lies in ensuring seamless interoperability between diverse, specialized simulation tools and heterogeneous data models. Standard co-simulation serves as the foundational integration layer, coupling specialized engines like CARLA (for high-fidelity sensing/physics), SUMO (for macroscopic flow), and ns-3 (for network protocol emulation). Pushing this boundary, differentiable co-simulation represents a paradigm shift, treating the entire multi-simulator pipeline as a single, end-to-end differentiable function. This architecture enables gradient-based optimization across the full stack, allowing system-wide parameters to be jointly tuned against global objectives, like from a vehicle's control logic in SV-DT using CARLA to traffic signal timing in ENV-DT using SUMO. Beyond execution coupling, semantic interoperability is essential for disparate DTs to interpret heterogeneous data sources correctly. This is achieved via Knowledge Graphs and Ontology-based Coordination [203], [204], which provide a formal, structured vocabulary defining domain entities (e.g., vehicles, infrastructure) and their complex interrelationships. To bridge organizational silos where proprietary schemas differ, Semantic Graph Alignment techniques [205] identify correspondences to merge disparate knowledge bases into a unified semantic layer. This shared understanding empowers ToTs ecosystem to query, reason over, and act upon shared information, enabling true global orchestration.

C. Cross-Domain Service Composition

1) Holistic Optimization and Knowledge Graph Embedding: Unlike conventional DTs that optimize isolated subsystems, ToTs framework offers a comprehensive, multimodal understanding of the vehicular environment. By capturing intricate correlations across vehicle control from SV-DT, inter-vehicle coordination from CV-DT, resource allocation from NET-DT, environmental dynamics from ENV-DT, and safety validation from SAFE-DT, ITCS enables the training and evaluation of algorithms against diverse, realistic scenarios, facilitating holistic system-level optimization. A pivotal advantage of this mature ecosystem is *Cross-Domain Service Composition*, where capabilities from traditionally siloed domains (mobility, communications, energy, and public safety) are integrated to synthesize novel, high-value services. Acting as the central orchestration platform, ToTs models, simulates, and optimizes the complex interactions between these disparate fields. To overcome the challenge of semantic interoperability, ToTs framework employs Knowledge Graph Embedding [206]. By mapping entities and relationships from multiple domains into a shared vector space, this technique allows ML models to reason across boundaries. For instance, quantifying how a specific communication protocol's latency profile directly impacts a vehicle's mechanical stability. A prime exemplar of this tight coupling is the joint optimization of control and communication. By simultaneously modeling vehicle dynamics and wireless link states, Multi-Timescale Deep Reinforcement Learning methods develop joint policies that execute intelligent trade-offs. These trade-offs are often mediated by a common metric such as the Value of Information, which explicitly quantifies the utility of specific data packets to the control task, thereby bridging the domains for holistic optimization as studied in [207].

2) Cross-Domain Transfer and Continual Learning: Particularly, facilitating effective service composition necessitates the seamless transfer of knowledge and predictive models across diverse operational domains, which is a challenge addressed by a suite of advanced machine learning paradigms. Cross-Domain Transfer Learning serves as the cornerstone, allowing models trained in a source domain (e.g., traffic prediction for a specific municipality) to be repurposed and fine-tuned for a target domain, thereby drastically reducing data acquisition and computational training costs. This capability branches into two strategic objectives: *Domain Adaptation*, which specializes a pre-trained model to excel in a novel environment, and *Domain Generalization*, which constructs a single, robust model capable of functioning effectively across heterogeneous domains without specific fine-tuning. To address highly dynamic environments where DT must evolve in real-time, Meta-Learning [208] is typically employed. This paradigm tunes the specific DT modules to rapidly adapt to new tasks or contexts with minimal, sparse data. Finally, to mitigate catastrophic forgetting during this evolution, Domain Continual Learning [209] ensures stable, cumulative knowledge acquisition, allowing DTs to retain mastery of prior domains while assimilating new operational realities throughout ITCS lifecycle.

3) LLM-Assisted Automated Service Composition: The culmination of this ToTs architecture is the automated composition and execution of cross-domain services, formally framed as a constraint-aware combinatorial optimization problem. To fulfill complex user requests, ITCS must dynamically select and orchestrate a specific set of atomic services and physical resources, such as sensing arrays, communication channels, or edge compute nodes, while strictly adhering to multi-domain constraints like network bandwidth, vehicle energy budgets, and traffic regulations. Operational examples include event-triggered navigation systems that fuse sensing services (incident detection) with dynamic re-planning algorithms [210], and resilient resource allocation frameworks that compose heterogeneous capabilities from terrestrial (RSU) and aerial (UAV) domains to ensure robust task offloading [211]. A frontier in automating this process is LLM–Assisted Composition [142], [212], [213]. In this paradigm, the LLM functions as a high-level cognitive orchestrator. By interpreting natural language requests, the model queries the underlying Knowledge Graph to discover available services and their interdependencies. It then formulates and solves the optimization problem, generating a valid, executable service composition plan on the fly.

IX. Open Challenges

A. Fidelity-Latency Trade-off

A fundamental and pervasive tension exists between the imperious need for high operational fidelity and the strict requirement for low-latency execution in modern DTs. High fidelity implies that the virtual model mirrors the physical ITCS with extreme precision, utilizing photorealistic visual rendering for camera sensor validation and complex physics calculations for vehicle dynamics. For instance, accurately simulating the friction coefficient of a tire on a wet road requires solving intricate differential equations and computational fluid dynamics models that consume significant computational resources. Similarly, generating synthetic image data that is indistinguishable from real-world video involves ray-tracing techniques that are computationally intensive. However, safety-critical transportation applications operate under tight timing constraints where vehicles traveling at highway speeds cover significant distances in mere milliseconds. Consequently, DT must process incoming sensor data and return actionable control guidance within this incredibly short window to remain relevant. If the simulation takes too long to process the physical state, the resulting advice arrives too late to prevent an accident, effectively rendering the insight useless. This creates a difficult engineering trade-off where increasing the realism of the model inevitably increases processing time and potential safety risks, forcing architects to choose between a slow but accurate truth and a fast but approximate estimate.

Current approaches [214]–[217] often rely on simplified proxy models to achieve real-time performance, but these linearized approximations frequently fail to capture the complex nonlinear behaviors seen in extreme scenarios, such as sudden hydroplaning or multi-vehicle collisions. While these reduced-order models suffice for nominal driving conditions, they lack the granularity required to validate safety in the rare, chaotic edge cases where DT is most needed. Furthermore, the hardware infrastructure complicates this balance significantly. Edge computing nodes located near roadside units offer the low network latency required for rapid feedback but typically lack the raw processing power to execute high-fidelity rendering or heavy physics engines. Conversely, centralized cloud data centers possess the necessary computational capacity to run detailed simulations but suffer from variable transmission delays and jitter that render them too slow for immediate control loops. Bridging this gap remains a significant open challenge, necessitating future research into efficient acceleration techniques, such as physics-informed machine learning surrogates [218], and hierarchical modeling [219] that can intelligently switch between high-fidelity simulation and fast, approximate inference depending on urgency of the situation.

B. Data Heterogeneity and Sparsity

The deployment of city-scale DTs faces a substantial barrier due to the chaotic and non-uniform nature of real-world data. Unlike the clean and structured datasets used in controlled experiments, data collected from physical vehicular networks is inherently heterogeneous. Different vehicles are equipped with sensors from various manufacturers, each possessing distinct resolution capabilities, calibration standards, and noise profiles. A modern autonomous vehicle might generate high-density LiDAR point clouds and 4K video streams, while an older connected vehicle might only broadcast basic GPS coordinates and speed telemetry. Integrating these highly disparate data sources into a unified DT requires sophisticated normalization pipelines that can mathematically reconcile inputs of varying quality. Without robust fusion techniques, a single noisy sensor from a low-quality source could corrupt the accuracy of the entire global model. Furthermore, this heterogeneity extends to the operational environment itself. Data distributions shift dramatically based on geography and weather. A model trained on telemetry from a sunny, flat coastal city will likely fail when applied to a mountainous region during a snowstorm because the underlying statistical properties of the friction and visibility are fundamentally different. This violation of the independent and identically distributed assumption, which underpins most machine learning algorithms, makes it difficult to build universal models that generalize across the entire fleet.

Equally critical is the challenge of data sparsity regarding safety-critical events. While ITCS generates petabytes of data, the vast majority of this information represents mundane and nominal driving scenarios, such as cruising on a highway or stopping at a red light. Dangerous events, such as a pedestrian suddenly stepping into traffic or a multi-vehicle collision, are statistically rare [220]. This phenomenon creates a severe data imbalance known as the long-tail problem. For instance, SAFE-DTs typically require thousands of examples to learn a specific pattern effectively. When the available dataset contains millions of hours of routine driving but only a few seconds of accident data, the resulting predictive models become heavily biased toward the normal state. Consequently, DT may perform exceptionally well at simulating everyday traffic flow but fail catastrophically when attempting to predict the outcome of an emergency situation. This scarcity creates a significant Sim-to-Real gap, where vehicle control agents trained in DT's simulation module struggle to adapt to the unpredictable edge cases of the physical road environments. Addressing this requires advanced techniques such as few-shot learning and the generation of synthetic adversarial data to artificially populate the sparse regions [221]–[223].

C. Security-Safety-Privacy Trilemma

Achieving a trustworthy ITCS-DT ecosystem requires simultaneously satisfying three critical objectives: system security, physical safety, and data privacy. However, these goals are frequently at odds, creating a complex trilemma where optimizing for one often degrades the performance of the others. For instance, maintaining different DT types relies on the broadcast of highly precise, unencrypted, and high-frequency telemetry data. To prevent a collision at an intersection, a vehicle must share its exact location, speed, and future trajectory with all nearby agents (e.g. for constructing a CV-DT). While this maximal information sharing ensures the high-fidelity of DT and resultant high-level of physical safety, it creates a severe privacy vulnerability. Malicious observers can easily eavesdrop on these open broadcasts to reconstruct a driver's detailed route history, infer their habits, and ultimately de-anonymize specific individuals. Conversely, prioritizing privacy by applying strict encryption or data minimization techniques hides this sensitive information but obscures the situational awareness required for DT mapping process and safe cooperative maneuvering.

This tension is particularly acute when applying standard privacy-preserving mechanisms like differential privacy [224]. These techniques protect individual data by injecting statistical noise into the reported values. In a typical data mining context, this noise is acceptable because it averages out over large datasets. However, in the context of an SV-DT controlling a physical vehicle, this added noise effectively corrupts the precise measurements needed for stability. If a vehicle reports its position with a privacy-preserving error margin of several meters, SV-DT cannot accurately calculate a safe braking distance, directly endangering physical safety. Similarly, robust security measures often introduce unacceptable latency overheads. Verifying complex cryptographic signatures and performing extensive integrity checks on every incoming message ensures that ITCS is not being spoofed by an attacker. Yet, these computational steps consume valuable milliseconds. In an emergency scenario where a vehicle must react instantly to a sudden obstacle, the time spent verifying the security credentials of a warning message could delay the actuation of the brakes, leading to a preventable accident. Therefore, the central challenge lies in developing lightweight, context-aware protocols that can dynamically balance these competing needs when deploying DTs in ITCS, enforcing strict security and privacy during normal operation while permitting immediate, transparent data access during safety-critical emergencies.

D. Energy Efficiency & Green Computing of ToTs

The deployment of city-scale DTs in ITCS introduces a profound paradox regarding environmental sustainability. While the primary objective of these systems is often to optimize urban efficiency, such as reducing traffic congestion to lower vehicle emissions or balancing smart grids to integrate renewable energy, the operation of DT itself consumes a staggering amount of electrical power. Maintaining a real-time, high-fidelity mirror of the complex physical transportation and communication systems requires the continuous operation of massive data centers and edge computing clusters on the roadside. These facilities must ingest petabytes of sensor streams, execute complex physics simulations, and retrain sophisticated ML models within DTs on a daily or even hourly basis. Consequently, the carbon footprint generated by the computational infrastructure can potentially negate the environmental benefits achieved through the system's optimization commands. This phenomenon challenges the fundamental premise of using DT solutions for green objectives, forcing researchers to scrutinize the net energy impact of the entire cyber-physical lifecycle. The energy burden is distributed across three primary stages: data transmission, DT maintenance, and real-time inference. First, transmitting high-bandwidth raw data, such as LiDAR point clouds and high-definition video, from thousands of vehicles to a central server places an immense load on the communication infrastructure. 5G/beyond base stations consume significant energy to support these high throughputs. Second, the ML models that power DTs, particularly DRL agents and large foundation models, require massive computational resources for training. Retraining these models to adapt to shifting traffic patterns involves running thousands of GPUs in parallel for extended periods, a process known to have a substantial carbon intensity. Third, the continuous execution of DTs requires servers to run simulation or prediction engines 24/7, regardless of actionable events.

Addressing this challenge requires a paradigm shift from performance-at-all-costs to energy-aware computing. Current research must move beyond simple efficiency metrics and adopt carbon-aware scheduling within ToTs, where non-urgent computational tasks, such as long-term model refinement in a specific DT (e.g., ENV-DT or SAFE-DT), are dynamically paused and resumed based on the real-time availability of green energy in the power grid. Techniques such as event-triggered processing [225], where the simulation only updates when a significant change occurs in the physical environment, and the use of neuromorphic hardware [226], which mimics the energy efficiency of the human brain, offer promising pathways to decouple the growth of digital intelligence from the growth of energy consumption.

E. Standardization & Semantic Interoperability

The realization of a comprehensive ToTs is currently obstructed by a pervasive lack of standardization across the disparate domains that must be integrated. The convergence of automotive engineering, urban infrastructure planning, and telecommunications creates a fragmented landscape where each sector adheres to its own legacy standards and proprietary protocols. For instance, the automotive industry typically utilizes high-precision formats like ASAM OpenDRIVE and OpenSCENARIO to define road geometry and dynamic behaviors for simulation. In contrast, urban planners and city managers rely on Geographic Information System standards such as CityGML or Shapefiles, which prioritize topological features and asset management over kinematic precision. Simultaneously, the telecommunications sector governs network connectivity through rigid 3GPP specifications that define quality of service parameters. This fundamental disconnect creates a digital Tower of Babel where an SV-DT cannot seamlessly interpret the map data provided by an ENV-DT, and a NET-DT cannot naturally parse the latency requirements of a safety application from a SAFE-DT. Consequently, ITCS integrators are forced to build bespoke, fragile translation layers and middleware wrappers for every new deployment, significantly increasing development costs and scalability.

Beyond the challenges of inconsistent file formats and communication protocols lies the deeper and more insidious problem of semantic interoperability. Syntactic compatibility ensures that two digital systems can exchange data strings, but it does not guarantee that they share a common understanding of what that data represents. A numeric value labeled speed transmitted by a vehicle could be interpreted by a traffic management system as kilometers per hour, while the sender intended it as meters per second. Such ambiguity, while trivial in human conversation, can lead to catastrophic failures in automated control loops. Furthermore, abstract concepts vary wildly between domains. The definition of an intersection in a navigation map might simply be a node where lines meet, whereas in a safety-critical autonomous driving simulation, it represents a complex zone of conflict rules, traffic light phases, and right-of-way logic. Without a unified semantic ontology that rigorously defines these entities and their relationships, the automated orchestration of cross-domain services from ToTs remains impossible. Future frameworks must move beyond simple data exchange APIs to adopt knowledge-graph-driven DT architectures where data is accompanied by machine-readable context, enabling disparate DTs to reason about shared information without ambiguity.

F. Hardware-Software Co-Design

The escalating computational demands of city-scale DTs are rapidly approaching the physical limits of conventional general-purpose hardware. Current implementations predominantly rely on standard Central Processing Units and Graphics Processing Units hosted in hyperscale cloud data centers. While these architectures excel at dense matrix operations common in deep learning, they are often ill-suited for the specific irregular workloads inherent to large-scale urban simulation. Twinning a metropolitan transportation network involves traversing massive, sparse graphs to calculate routing paths for millions of agents while simultaneously solving disparate differential equations for vehicle dynamics. These tasks create irregular memory access patterns that cause frequent cache misses on standard hardware, leading to a phenomenon known as the Von Neumann bottleneck. In this scenario, the processor spends more time waiting for data to be fetched from memory than actually performing computations, resulting in severe inefficiencies where a vast amount of theoretical teraflops remain unutilized.

To overcome this silicon wall, the field must embrace a paradigm of hardware-software co-design, where the physical processor and the logic of DT software are developed in unison. This transition necessitates the move away from monolithic, one-size-fits-all processors toward Domain-Specific Architectures tailored explicitly for DT primitives. For example, Field-Programmable Gate Arrays [227] can be customized to execute specific physics solvers or collision detection pipelines directly in hardware circuits, offering orders of magnitude lower latency than software running on a general-purpose instruction set. Similarly, emerging Graph Processing Units are designed to handle the massive connectivity of traffic networks natively, enabling DTs to update the state of millions of connected entities in parallel without the data movement overheads of traditional chips. Furthermore, this co-design extends to the software stack itself. Legacy co-simulation engines in ToTs written for standard x86 architectures cannot simply be ported to these specialized accelerators. They must be fundamentally re-architected to expose parallelism and data locality that matches the underlying hardware fabric. Future DT platforms will likely resemble heterogeneous compute clusters, where a central orchestrator dispatches prediction and simulation tasks to tensor cores, physics tasks to FPGA accelerators, and graph traversal across DTs to graph processors. Mastering this complex interplay between specialized silicon and optimized code is the only viable path to achieving the real-time, city-scale fidelity required for the next generation of ITCS-DT development.

X. Future Research Directions

This section outlines future research directions and potential solutions to the challenges discussed in Sec. IX. The conceptual relationships among them are illustrated in Fig. 8.

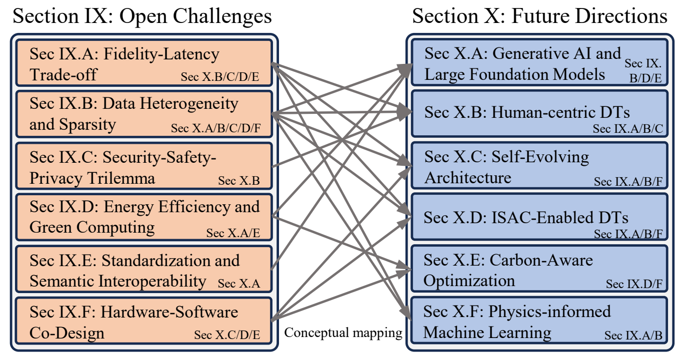

Fig. 8: Conceptual mapping between the Open Challenges and the Future Directions.

A. *Generative AI & Foundation Models for DT construction*

The integration of Generative AI (GenAI) and large foundation models can shift DTs from static, replay-based replicas to *generative safety testbeds* that proactively synthesize rare yet high-risk situations. This direction directly addresses the *Data Heterogeneity and Sparsity* challenge in Section IX-B: safety-critical events are statistically rare, yielding a long-tail distribution where many hazardous conditions are missing from logs. Instead of waiting years to observe a specific failure-prone circumstance, a DT can learn a conditional scenario model $p_\theta(\mathbf{x} \mid \mathbf{c})$ and sample thousands of physically plausible variants. In practice, a conditional diffusion generator can be trained with

$$\min_{\theta} \ \mathbb{E}_{t,\mathbf{x},\boldsymbol{\epsilon}} \left[\|\boldsymbol{\epsilon} - \epsilon_\theta(\mathbf{x}_t, t, \mathbf{c})\|_2^2 \right], \tag{1}$$

where $\mathbf{x}$ denotes a compact scenario representation (e.g., multi-agent states and local geometry), $\mathbf{c}$ is DT-extracted condition vector, and $\mathbf{x}_t$ is the noised state at diffusion step t with injected noise $\boldsymbol{\epsilon}$. A practical choice is $\mathbf{c} = [\mathbf{m}; \boldsymbol{\omega}; \mathbf{s}; \boldsymbol{\iota}]$, concatenating map semantics $\mathbf{m}$, environment/friction state $\boldsymbol{\omega}$, surrounding traffic context $\mathbf{s}$, and test intent $\boldsymbol{\iota}$. Sampled scenarios are then filtered by lightweight feasibility checks (e.g., bounded acceleration/jerk, collision-free geometry, and rule-consistent trajectories) before being added to a DT scenario bank for controller evaluation, optionally with coverage-aware sampling to prioritize under-tested regions of the ODD.

Beyond scenario generation, LLMs can mitigate the *Standardization & Semantic Interoperability* bottleneck in Section IX-E. Today, fragmented proprietary formats often require brittle hand-crafted translators across sensors, maps, simulators, and infrastructure APIs. An LLM can serve as a semantic mediator that parses unstructured specifications and heterogeneous documentation to produce schema mappings, unit normalization, and executable glue code, enabling an LLM-as-Orchestrator workflow that converts natural-language intent into validated DT invocations. However, deploying massive models introduces new tension with the energy and latency constraints in Section IX-D. Future work should therefore prioritize efficient DT-ready pipelines–distillation from a large teacher DT into a lightweight student DT, quantization and caching for low-latency inference, and hybrid execution where heavy generative reasoning runs off-board while safety-critical loops rely on compact, verifiable models on edge hardware.

B. *Human-centric DTs (HC-DTs)*

While current SV-DTs and ENV-DTs excel at modeling vehicle kinematics and infrastructure dynamics, they often treat humans—drivers, pedestrians, and cyclists—as simplified, deterministic agents. A critical frontier is the development of Human-centric DTs in ITCS that explicitly model cognition, intent, and affect, so that safety assessment reflects how real incidents unfold. In practice, traffic risk is rarely dictated solely by braking distance or friction; it is frequently amplified by latent human states such as fatigue, distraction, stress, or risk-seeking behavior. To bridge the *Data Heterogeneity and Sparsity gap* in Section IX-B, future DTs must move beyond tracking *where* a human is to inferring *why* they behave a certain way and *how* their behavior may evolve under changing context. This can be operationalized by treating human factors as partially observed dynamics: a DT fuses multimodal observations (e.g., gaze/pose, heart-rate variability, steering micro-corrections, headway fluctuations, smartphone interaction cues) into a latent cognitive state and uses it to anticipate erratic maneuvers before they physically manifest.

A practical formulation is to represent the human cognitive state as a discrete (or low-dimensional continuous) latent variable $z_t \in \mathcal{Z}$ inferred from multimodal observations o_t. DT maintains an online belief $b_t(z) \triangleq p(z_t = z \mid o_{1:t})$ and updates it with a Bayesian filter:

$$b_t(z) = \eta \, p(o_t|z) \sum_{z'} p(z|z') \, b_{t-1}(z'), \tag{2}$$

where $p(o_t|z)$ is the observation likelihood, $p(z|z')$ is the latent-state transition model, and η normalizes b_t to a probability distribution. The belief can be mapped to a risk score $r_t = g(b_t)$ (e.g., expected hazard probability or expected policy deviation), which then triggers adaptive DT behaviors such as switching to conservative planning, tightening safety buffers, re-weighting prediction uncertainty, or issuing context-aware alerts. To respect the Fidelity-Latency Trade-off in Section IX-A, such human models should favor lightweight inference (e.g., small-state $\mathcal{Z}$, amortized likelihoods, or low-rank embeddings) and modular "theory-of-mind" components that approximate intent and compliance in real time without full psychological simulation.

However, deep human sensing required for human-centric DTs immediately intensifies the *Security-Safety-Privacy Trilemma* in Section IX-C. A DT that models physiological and psychological state poses a privacy risk beyond location tracking, since compromise could reveal sensitive attributes. Consequently, human-centric DTs should adopt privacy-preserving computation by design: raw biometric streams remain on-device, and learning is performed via federated optimization. In this setting, only aggregated updates are shared, which can be further strengthened with secure aggregation and/or encrypted computation to reduce leakage. Overall, the long-term goal is a DT that is not only a physics-consistent controller, but also a human-aware co-pilot that adapts its assistance and safety envelope to the inferred human condition under strict privacy constraints.

C. Self-Evolving ToTs Architecture

Current ToTs implementation rationale is predominantly static artifacts: once deployed and calibrated, their structural logic remains fixed until a human engineer manually updates the stack. This rigidity becomes a vulnerability under the *Data Heterogeneity and Sparsity* challenge in Section IX-B, because the statistical properties of traffic, sensing, and infrastructure can shift dramatically across regions, seasons, and operational policies. A model optimized for one domain (e.g., dense urban arterials) may degrade sharply when transferred to another (e.g., suburban corridors), not merely due to parameter drift but because the *appropriate model complexity and module composition* changes. Future research should therefore pivot toward a self-evolving architecture: ToTs that can autonomously adapt both their parameters and their computational structure, using mechanisms such as modular routing, dynamic pruning/expansion, and AutoML/NAS-driven topology updates.

A practical self-evolving ToTs can be framed as a budget-aware architecture controller that selects a model configuration $m \in \mathcal{M}$ of its contained DT (e.g., lightweight kernel, full perception, weather-aware fusion) to minimize safety risk under latency and energy constraints. This formulation directly supports the *Fidelity-Latency Trade-off* in Section IX-A: e.g., during routine conditions, ENV-DT component can route to a compact configuration that reduces compute, while in complex geometry, occlusions, or adverse weather it can activate higher-fidelity modules (e.g., heavier perception backbones, uncertainty-aware predictors, or multi-sensor fusion) to preserve safety margins. Importantly, this adaptation is not a heuristic toggle; it is an explicit constrained selection over a discrete set of deployable configurations with measurable budgets. Beyond short-term elastic fidelity, ToTs can improve their *baseline* architecture through a periodic AutoML/NAS loop that searches for a better topology as new data accumulates. A standard bilevel view is

$$\min_{\alpha} \; \mathcal{L}_{val}(w^*(\alpha), \alpha) \, , w^*(\alpha) = \arg\min_{w} \; \mathcal{L}_{train}(w, \alpha), \tag{3}$$

where α encodes architectural choices (e.g., depth/width, attention or temporal blocks, sensor-fusion operators, and module connectivity) and w are the corresponding weights for different twin models. The resulting candidate is promoted only if it satisfies the operational constraints and demonstrates robustness across shifted conditions (e.g., different corridors, seasons, and sensor health states). Together, the online controller and offline search yield ToTs that can right-size themselves to each context in real time while continuously upgrading their structural capacity over longer horizons, reducing dependence on manual redesign and improving resilience to distribution shift.

D. ISAC-Enabled NET-DTs

The evolution from 5G toward 6G enables a structural shift in which DTs move from Over-The-Top cloud applications to *network-native* services co-designed with the communication fabric. In current deployments, the network primarily functions as a transport layer that forwards bits with limited awareness of the time-critical semantics carried by system state updates. A 6G-native NET-DT instead leverages *In-Network Computing* [228], where parts of the twinning pipeline (state estimation, lightweight prediction, and fast synchronization) are executed within MEC or even base-station-side compute. This directly mitigates the *Fidelity-Latency Trade-off* in Section IX-A: by pushing the twin update loop to the extreme edge, the system can reduce control-loop delay and maintain tighter alignment between physical processes and their digital counterparts. A convenient abstraction is the end-to-end latency budget and to minimize it through edge execution, caching of twin fragments, and fast multicast/broadcast primitives, enabling near real-time synchronization for safety-critical loops and emerging interactive applications.

Beyond low-latency compute, 6G is expected to operationalize ISAC (Integrated Sensing and Communication), turning the network infrastructure into a pervasive sensing substrate. This capability complements vehicle-mounted sensing and helps alleviate the *Data Heterogeneity and Sparsity* challenge in Section IX-B, since ISAC provides a standardized perspective that is less sensitive to differences in individual vehicle sensor suites. In an ISAC-enabled NET-DT, the communication system produces measurements (e.g., range, Doppler, and angle-related features) that relate to the traffic and environment states, which can be fused with onboard perception to update NET-DT state online (e.g., via Kalman-type filtering), and the updated state can then be broadcast back to vehicles and applications, closing a low-latency *sensing–compute–sync* loop. Consequently, NET-DT will become a unified city-scale fabric where the boundary between communication and sensing blurs, supporting high-resolution situational awareness even in regions with sparse connected vehicles, degraded visibility, or inconsistent sensor quality.

E. Carbon-Aware Optimization of City-Scale DTs

As highlighted in Section IX-D, the growing compute demand of city-scale DTs creates an environmental paradox: higher-fidelity intelligence often implies higher electricity consumption. Future frameworks should therefore evolve from being merely energy-efficient to being explicitly carbon-aware, treating emissions as a first-class objective rather than an after-the-fact metric. In this paradigm, DT is not only a mirror of ITCS but also a sustainability orchestrator that co-optimizes digital workloads with the carbon dynamics of the power grid. By integrating real-time (or forecast) grid signals such as renewable generation, marginal emission factors, and time-varying carbon intensity, DT can shift non-urgent tasks (e.g., offline retraining, batch data curation, large-scale simulation sweeps, or scenario generation) away from high-carbon periods and concentrate them when the grid mix is cleaner. Formally, if $g(t)$ denotes grid carbon intensity and $P(t)$ denotes DT computing power, a carbon-aware scheduler can minimize time-integrated emissions via

$$\min_{P(t)} \int_{t_0}^{t_1} g(t)\, P(t)\, dt \quad \text{s.t.} \quad \text{QoS constraints}, \tag{4}$$

thereby decoupling the growth of DT intelligence from the worst-case carbon footprint of its underlying infrastructure. In practice, this can be implemented as a job-level controller that assigns start times and power caps to workloads under deadlines, while respecting operational constraints such as service-level objectives for latency-sensitive DT functions.

Beyond optimizing the individual twin's footprint, ToTs need to reduce emissions in the physical transportation-energy system itself by coupling CV-DT (for vehicle mobility) with an ENV-DT (for infrastructure-level energy consumption). This cross-domain coupling supports carbon-aware decision making such as green routing and smart charging, where electric vehicles are guided not only by travel time but also by the marginal carbon intensity of charging opportunities along candidate paths. At a system level, ENV-DT-driven coordination can further enable vehicle-to-grid (V2G) dispatch by forecasting demand spikes and scheduling distributed battery discharge during critical shortages, improving renewable utilization and reducing peaker-plant reliance. By embedding these functionalities into ITCS-DT control loop, carbon-aware optimization becomes measurable and enforceable, enabling DT deployments that are both computationally scalable and environmentally defensible.

F. Physics-informed DTs in ITCS

A persistent dichotomy in DT modeling for ITCS is the trade-off between white-box physics-based solvers and black-box data-driven models. Physics-based simulators (e.g., vehicle dynamics models, traffic-flow partial differential equations (PDEs), such as computational fluid dynamics solvers for aerodynamic analysis) provide interpretability and strong extrapolation guarantees. However, they are often computationally prohibitive for real-time vehicular control, cooperative perception, and network adaptation under the *Fidelity-Latency Trade-off* discussed in Section IX-A. In contrast, pure deep learning models enable fast inference for trajectory prediction, channel estimation, or traffic-state forecasting, yet they typically require large-scale labeled datasets and may violate fundamental physical or safety constraints when operating out of distribution. Physics-Informed Machine Learning (PIML) offers a principled middle ground by embedding governing transportation and communication laws directly into the learning process. These laws may include vehicle kinematics and dynamics constraints, traffic conservation equations, wireless propagation models, or queue evolution dynamics. Empirical studies [229], [230] in physics-informed surrogate modeling report one to three orders-of-magnitude acceleration compared to full numerical solvers, reducing simulation times from minutes or seconds per scenario to millisecond-level inference while maintaining low residual error. In ITCS context, this acceleration enables near-real-time vehicular DT updates for trajectory forecasting, traffic-state evolution, or channel prediction, thereby directly mitigating the *Fidelity–Latency Trade-off*. Particularly, the complexity shift from iterative PDEs solving to feed-forward evaluation makes PIML particularly attractive for edge-deployed NET-DT or SV-DT components, where computational budgets are constrained yet strict latency guarantees are required. Concretely, a PIML-enabled DT learns a neural surrogate $u_\theta(\cdot)$ that jointly fits observed data while penalizing violations of physical or network constraints:

$$\mathcal{L}(\theta) = \mathcal{L}_{\mathrm{D}}(\theta) + \lambda\, \mathcal{L}_{\mathrm{P}}(\theta) + \mu\, \mathcal{L}_{\mathrm{B}}(\theta), \tag{5}$$

where $\mathcal{L}_{\mathrm{D}}$ measures data fidelity such as trajectory reconstruction error, channel prediction error, or traffic-state mismatch, $\mathcal{L}_{\mathrm{B}}$ enforces boundary or initial conditions such as road topology constraints, signal timing plans, or communication protocol limits, and $\mathcal{L}_{\mathrm{P}}$ penalizes violations of governing laws. The hyperparameters λ and μ balance empirical accuracy against physical consistency [231]. This hybrid modeling paradigm directly addresses the *Data Heterogeneity and Sparsity challenge* in Section IX-B. In ITCS, rare events such as emergency braking, abrupt lane changes, congestion shockwaves, link blockage at mmWave bands, or network overload are typically under-sampled. Unconstrained DTs may therefore generate physically implausible trajectories, unsafe maneuvers, or unstable network-state predictions. PIML constrains the hypothesis space to physically admissible solutions. For example, even with limited supervision for a specific vehicular scenario, a SV-DT is discouraged from predicting non-feasible accelerations, instantaneous velocity discontinuities, negative traffic densities, or queue dynamics that violate flow conservation. Formally, the governing transportation or communication law is represented by an operator $\mathcal{F}(\cdot)$, such as a traffic-flow PDE residual, a discrete conservation constraint on vehicle counts, a vehicular kinematic model, or a wireless channel evolution equation. The physics regularization term is defined as $\mathcal{L}_{\mathrm{P}}(\theta) = \|\mathcal{F}(u_\theta)\|_2^2$, where $\mathcal{F}(u_\theta)$ can be computed via automatic differentiation for continuous formulations or through discrete operators consistent with DT implementation loop. Anchoring the learning process to invariant physical structure reduces label demand and improves robustness under distribution shift, since predictions remain consistent with transportation dynamics and network constraints rather than relying solely on empirical correlations [232].

XI. Conclusion

This survey presents a comprehensive overview of DTs in connected intelligent transportation and communication systems, marking the architectural evolution from isolated simulations to a holistic ToTs ecosystem. By integrating vehicular dynamics, communication networks, and energy infrastructures, this system-of-systems approach enables unprecedented cross-domain interoperability and service composition. We have analyzed critical enablers, including high-fidelity co-simulation, semantic knowledge graphs, and distributed intelligence, while identifying persistent challenges. Ultimately, bridging these gaps through emerging technologies like Generative AI, physics-informed learning, and 6G-native computing is essential to realizing a truly autonomous, sustainable, and safety-critical transportation future.

XII. Acknowledgment

The authors would like to thank Dr. Dusit Niyato for valuable discussions and helpful suggestions on improving the writing and presentation of this paper.

References

[1] T. Gong, L. Zhu, F. R. Yu, and T. Tang, "Edge intelligence in intelligent transportation systems: A survey," *IEEE TITS*, 2023.

[2] X. Deng, L. Wang, J. Gui, P. Jiang, X. Chen, F. Zeng, and S. Wan, "A review of 6g autonomous intelligent transportation systems: Mechanisms, applications and challenges," *JSA*, 2023.

[3] X. Li, J. Wei, H. Wang, L. Dong, R. Chen, C. Yi, J. Cai, D. Niyato, and X. Shen, "Towards intelligent transportation with pedestrians and vehicles in-the-loop: A surveillance video-assisted federated digital twin framework," *IEEE Network*, 2025.

[4] S. Alsubai, A. Alqahtani, A. Alanazi, and M. Bhatia, "Digital-twin-inspired iot-assisted intelligent performance analysis framework for electric vehicles," *IEEE IoT-J*, 2024.

[5] I. Saba, M. Ullah, and M. Tariq, "Advancing electric vehicle battery analysis with digital twins in intelligent transportation systems," *IEEE TITS*, 2024.

[6] Z. Zhang, Z. Peng, H. Yu, M. Chen, and Y. Liu, "Digital network twins for next-generation wireless: Creation, optimization, and challenges," *arXiv:2410.18002*, 2024.

[7] Z. Zhang, M. Fang, D. Chen, X. Yang, and Y. Liu, "Synergizing ai and digital twins for next-generation network optimization, forecasting, and security," *IEEE Wireless Communications*, 2025.

[8] X. Hu, S. Li, T. Huang, B. Tang, R. Huai, and L. Chen, "How simulation helps autonomous driving: A survey of sim2real, digital twins, and parallel intelligence," *IEEE TIV*, 2023.

[9] Z. Wang, R. Gupta, K. Han, H. Wang, A. Ganlath, N. Ammar, and P. Tiwari, "Mobility digital twin: Concept, architecture, case study, and future challenges," *IEEE IoT-J*, 2022.

[10] Z. Zhang, Y. Liu, Z. Peng, M. Chen, D. Xu, and S. Cui, "Digital twin-assisted data-driven optimization for reliable edge caching in wireless networks," *IEEE JSAC*, 2024.

[11] B. Hazarika, K. Singh, A. Paul, and T. Q. Duong, "Hybrid machine learning approach for resource allocation of digital twin in uav-aided internet-of-vehicles networks," *IEEE TIV*, 2023.

[12] G. Xu, J. Chen, Z. Wang, A. Zhou, M. Schrader, J. Bittle, and Y. Shao, "Enhancing traffic safety analysis with digital twin technology: Integrating vehicle dynamics and environmental factors into microscopic traffic simulation," *arXiv:2502.09561*, 2025.

[13] X. Wang, M. Hao, M. Wu, C. Shang, R. Yu, J. Kang, Z. Xiong, and Y. Wu, "Digital twin-assisted safety control for connected automated vehicles in mixed-autonomy traffics," *IEEE IoT-J*, 2024.

[14] J. Duan, Z. Wang, and X. Jing, "Digital twin test method with lte-v2x for autonomous vehicle safety test," *IEEE IoT-J*, 2024.

[15] O. Zheng, M. Abdel-Aty, L. Yue, A. Abdelraouf, Z. Wang, and N. Mahmoud, "Citysim: A drone-based vehicle trajectory dataset for safety-oriented research and digital twins," *TRR*, 2024.

[16] M. Elloumi, M. Z. Hassan, and G. Kaddoum, "Spectrum sharing in internet-of-vehicles networks: Digital twin-empowered proactive interference management approach," *IEEE TNSM*, 2025.

[17] D. Chen, Y. Gong, and X. Yang, "Advanced longitudinal control and collision avoidance for high-risk edge cases in autonomous driving," *arXiv preprint arXiv:2504.18931*, 2025.

[18] J. Liu, X. Sheng, L. Tan, W. Zhang, P. Zhang, and K. Liu, "Dpdrf: Dynamic predictive driving risk field based on multi-agent trajectory prediction and digital twins system," *IEEE TVT*, 2025.

[19] Y. Liang, Z. Yin, L. Nie, and Y. Ba, "Shared steering control with predictive risk field enabled by digital twin," *IEEE TIV*, 2023.

[20] "Digital twin network – requirements and architecture," International Telecommunication Union, Y. 3090 Recommendations, 2024.

[21] "Digital twin network - capability levels and evaluation methods," International Telecomm. Union, Y. 3091 Recommendations, 2024.

[22] ISO Central Secretary, "ISO 23247-1:2021 Automation systems and integration — Digital twin framework for manufacturing — Part 1: Overview and general principles," International Organization for Standardization, Geneva, Switzerland, Standard, 2021.

[23] 3rd Generation Partnership Project (3GPP), "Study on Network Controlled Interactive Services," 3GPP, Technical Report, 2019.

[24] ASAM e.V., *ASAM OpenDRIVE: File Format for the Description of Road Networks*, 1st ed., 2021. [Online]. Available: https://www.asam.net/standards/detail/opendrive/

[25] G. Bhatti, H. Mohan, and R. R. Singh, "Towards the future of smart electric vehicles: Digital twin technology," *Renewable and Sustainable Energy Reviews*, 2021.

[26] T. Deng, K. Zhang, and Z.-J. M. Shen, "A systematic review of a digital twin city: A new pattern of urban governance toward smart cities," *JMSE*, 2021.

[27] Z. Hu, S. Lou, Y. Xing, X. Wang, D. Cao, and C. Lv, "Review and perspectives on driver digital twin and its enabling technologies for intelligent vehicles," *IEEE TIV*, 2022.

[28] M. S. Irfan, S. Dasgupta, and M. Rahman, "Toward transportation digital twin systems for traffic safety and mobility: A review," *IEEE IoT-J*, 2024.

[29] J. Guo, M. Bilal, Y. Qiu, C. Qian, X. Xu, and K.-K. R. Choo, "Survey on digital twins for internet of vehicles: Fundamentals, challenges, and opportunities," *DCN*, 2024.

[30] Y. Pan, L. Lei, G. Shen, X. Zhang, and P. Cao, "A survey on digital twin networks: Architecture, technologies, applications and open issues," *IEEE IoT-J*, 2025.

[31] B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter, "Differentiable mpc for end-to-end planning and control," *Neurips*, 2018.

[32] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *ICML*, 2017.

[33] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska *et al.*, "Overcoming catastrophic forgetting in neural networks," *Proceedings of the national academy of sciences*, 2017.

[34] Z. Lv, J. Guo, A. K. Singh, and H. Lv, "Digital twins based vr simulation for accident prevention of intelligent vehicle," *IEEE TVT*, 2022.

[35] A. Stocco, M. Weiss, M. Calzana, and P. Tonella, "Misbehaviour prediction for autonomous driving systems," in *ACM/IEEE ICSE*, 2020.

[36] J. Li, W. Zhan, Y. Hu, and M. Tomizuka, "Generic tracking and probabilistic prediction framework and its application in autonomous driving," *IEEE TITS*, 2019.

[37] W. Hu, Y. He, Z. Liu, J. Tan, M. Yang, and J. Chen, "Toward a digital twin: time series prediction based on a hybrid ensemble empirical mode decomposition and bo-lstm neural networks," *Journal of Mechanical Design*, 2021.

[38] K. S. S. Alamin, Y. Chen, E. Macii, M. Poncino, and S. Vinco, "Advancing electric vehicle battery management: A data-driven digital twin approach for real-time monitoring and performance enhancement," *IEEE TVT*, 2025.

[39] X. Liao, X. Zhao, Z. Wang, Z. Zhao, K. Han, R. Gupta, M. J. Barth, and G. Wu, "Driver digital twin for online prediction of personalized lane-change behavior," *IEEE IoT-J*, 2023.

[40] Y. Hu, M. Wu, J. Kang, and R. Yu, "D-tracking: digital twin enabled trajectory tracking system of autonomous vehicles," *IEEE TVT*, 2024.

[41] L. Bartolucci, E. Cennamo, S. Cordiner, V. Mulone, F. Pasqualini, and M. A. Boot, "Digital twin of a hydrogen fuel cell hybrid electric vehicle: Effect of the control strategy on energy efficiency," *International journal of hydrogen energy*, 2023.

[42] D. Liu, W. Wang, L. Wang, H. Jia, and M. Shi, "Dynamic pricing strategy of electric vehicle aggregators based on ddpg reinforcement learning algorithm," *IEEE access*, 2021.

[43] Y. Chang and X. Wang, "Vehicle trajectory prediction with multimodal and dynamics-aware interaction neural networks," *IEEE TVT*, 2024.

[44] H. Shao, L. Wang, R. Chen, H. Li, and Y. Liu, "Safety-enhanced autonomous driving using interpretable sensor fusion transformer," in *Conference on Robot Learning*, 2023.

[45] D. J. Yeong, K. Panduru, and J. Walsh, "Exploring the unseen: A survey of multi-sensor fusion and the role of explainable ai (xai) in autonomous vehicles," *Sensors*, 2025.

[46] T. Nyberg, C. Pek, L. Dal Col, C. Norén, and J. Tumova, "Risk-aware motion planning for autonomous vehicles with safety specifications," in *IEEE IV*, 2021.

[47] X. Liu, Z. Li, L. Lei, G. Shen, S. Cai, and X. Liu, "Dynamic data-driven digital twin network construction and calibration for uav swarms," *IEEE IoT-J*, 2025.

[48] H. He, Y. Zhang, R. Xiong, and C. Wang, "A novel gaussian model based battery state estimation approach: State-of-energy," *Applied Energy*, 2015.

[49] J. Li, S. Guo, W. Liang, J. Wang, Q. Chen, Y. Zeng, B. Ye, and X. Jia, "Digital twin-enabled service provisioning in edge computing via continual learning," *IEEE TMC*, 2023.

[50] C. Zhang, Q. Zhou, B. Shuai, H. Williams, Y. Li, L. Hua, and H. Xu, "Dedicated adaptive particle swarm optimization algorithm for digital twin based control optimization of the plug-in hybrid vehicle," *IEEE TTE*, 2022.

[51] D. Chen, Y. Gong, and X. Yang, "Deep reinforcement learning for advanced longitudinal control and collision avoidance in high-risk driving scenarios," *arXiv:2404.19087*, 2024.

[52] H. Du, S. Leng, J. He, K. Xiong, and L. Zhou, "Digital twin empowered cooperative trajectory planning of platoon vehicles for collision avoidance with unexpected obstacles," *DCN*, 2024.

[53] Z. Xun, S. Zhang, Y. Xun, D. Mu, B. Mao, and H. Guo, "An adaptive vehicle trajectory prediction scheme based on digital twin platform," *IEEE IoT-J*, 2025.

[54] K. Kušić, R. Schumann, and E. Ivanjko, "A digital twin in transportation: Real-time synergy of traffic data streams and simulation for virtualizing motorway dynamics," *AEI*, 2023.

[55] B. Fan, Z. Su, Y. Chen, Y. Wu, C. Xu, and T. Q. Quek, "Ubiquitous control over heterogeneous vehicles: A digital twin empowered edge ai approach," *IEEE Wireless Communications*, 2022.

[56] H. Wang, C. Fennell, S. Poojary, and M. Liu, "A sumo-based digital twin for evaluation of conventional and electric vehicle networks," *arXiv:2507.10280*, 2025.

[57] C. Schwarz and Z. Wang, "The role of digital twins in connected and automated vehicles," *IEEE ITSM*, 2022.

[58] R. Song, C. Liang, H. Cao, Z. Yan, W. Zimmer, M. Gross, A. Festag, and A. Knoll, "Collaborative semantic occupancy prediction with hybrid feature fusion in connected automated vehicles," in *CVF CVPR*, 2024.

[59] S. Steyer, G. Tanzmeister, and D. Wollherr, "Grid-based environment estimation using evidential mapping and particle tracking," *IEEE TIV*, 2018.

[60] K. Wang, C. Shen, X. Li, and J. Lu, "Uncertainty quantification for safe and reliable autonomous vehicles: A review of methods and applications," *IEEE TITS*, 2025.

[61] Z. Yang, X. Zhang, D. Wu, R. Wang, P. Zhang, and Y. Wu, "Efficient asynchronous federated learning research in the internet of vehicles," *IEEE IoT-J*, 2022.

[62] Y. Li, X. Xu, G. Huang, M. Yao, L. Sun, and J. Xu, "Vsfl: Trajectory prediction framework based on validity-aware semi-asynchronous federated learning in internet of vehicles," *Computer Communications*, 2024.

[63] Y. Xing, C. Lv, and D. Cao, "Personalized vehicle trajectory prediction based on joint time-series modeling for connected vehicles," *IEEE TVT*, 2019.

[64] W.-B. Kou, Q. Lin, M. Tang, S. Xu, R. Ye, Y. Leng, S. Wang, G. Li, Z. Chen, G. Zhu *et al.*, "pfedlvm: A large vision model (lvm)-driven and latent feature-based personalized federated learning framework in autonomous driving," *IEEE TITS*, 2025.

[65] Q. Zhang, K. Yu, Z. Guo, S. Garg, J. J. Rodrigues, M. M. Hassan, and M. Guizani, "Graph neural network-driven traffic forecasting for the connected internet of vehicles," *IEEE TNSE*, 2021.

[66] G. Liu, Y. Xu, Z. He, Y. Rao, J. Xia, and L. Fan, "Deep learning-based channel prediction for edge computing networks toward intelligent connected vehicles," *IEEE Access*, 2019.

[67] P. Arthurs, L. Gillam, P. Krause, N. Wang, K. Halder, and A. Mouzakitis, "A taxonomy and survey of edge cloud computing for intelligent transportation systems and connected vehicles," *IEEE TITS*, 2021.

[68] A. Mallik, D. Chen, K. Han, J. Xie, and Z. Han, "Unleashing the true power of age-of-information: Service aggregation in connected and autonomous vehicles," in *IEEE ICC*, 2024.

[69] G. D'Angelo, A. Castiglione, and F. Palmieri, "A cluster-based multidimensional approach for detecting attacks on connected vehicles," *IEEE IoT-J*, 2020.

[70] W. Wang, A. Ramesh, J. Zhu, J. Li, and D. Zhao, "Clustering of driving encounter scenarios using connected vehicle trajectories," *IEEE TIV*, 2020.

[71] M. Aloqaily, S. Otoum, I. Al Ridhawi, and Y. Jararweh, "An intrusion detection system for connected vehicles in smart cities," *Ad Hoc Networks*, 2019.

[72] M. Y. Ansari, A. Ahmad, S. S. Khan, G. Bhushan, and Mainuddin, "Spatiotemporal clustering: a review," *Artificial Intelligence Review*, 2020.

[73] G. J. Oyewole and G. A. Thopil, "Data clustering: application and trends," *Artificial intelligence review*, 2023.

[74] M. Yang, B. Ai, R. He, G. Wang, L. Chen, X. Li, C. Huang, Z. Ma, Z. Zhong, J. Wang *et al.*, "Measurements and cluster-based modeling of vehicle-to-vehicle channels with large vehicle obstructions," *IEEE TWC*, 2020.

[75] B. Fan, Y. Wu, Z. He, Y. Chen, T. Q. Quek, and C.-Z. Xu, "Digital twin empowered mobile edge computing for intelligent vehicular lane-changing," *IEEE Network*, 2021.

[76] H. Liu, Z. Huang, X. Mo, and C. Lv, "Augmenting reinforcement learning with transformer-based scene representation learning for decision-making of autonomous driving," *IEEE TITS*, 2024.

[77] A. M. Aslam, R. Chaudhary, and A. Bhardwaj, "An aiot-enabled digital twin cavs with a drl-based framework for trajectory planning," *IEEE TITS*, 2025.

[78] X. Liao, S. Leng, Y. Sun, K. Zhang, and M. A. Imran, "A digital twin-based traffic guidance scheme for autonomous driving," *IEEE IoT-J*, 2024.

[79] J. Wang, C. Zhang, Z. Yang, M. Dang, P. Gao, and Y. Feng, "Research on digital twin vehicle stability monitoring system based on side slip angle," *IEEE TITS*, 2023.

[80] S. Venkatesan, K. Manickavasagam, N. Tengenkai, and N. Vijayalakshmi, "Health monitoring and prognosis of electric vehicle motor using intelligent-digital twin," *IET Electric Power Applications*, 2019.

[81] X. Zhou, X. Zheng, X. Cui, J. Shi, W. Liang, Z. Yan, L. T. Yang, S. Shimizu, and K. I.-K. Wang, "Digital twin enhanced federated reinforcement learning with lightweight knowledge distillation in mobile networks," *IEEE JSAC*, 2023.

[82] Z. Wang, K. Han, and P. Tiwari, "Digital twin-assisted cooperative driving at non-signalized intersections," *IEEE TIV*, 2021.

[83] L. Chen and C. Englund, "Cooperative intersection management: A survey," *IEEE TITS*, 2015.

[84] R. Chen, S. Han, J. Xu, and H. Su, "Point-based multi-view stereo network," in *CVF ICCV*, 2019.

[85] Y. Yao, Z. Luo, S. Li, T. Shen, T. Fang, and L. Quan, "Recurrent mvsnet for high-resolution multi-view stereo depth inference," in *CVF CVPR*, 2019.

[86] J. Guan, X. Yang, V. C. Lee, W. Liu, Y. Li, L. Ding, and B. Hui, "Full field-of-view pavement stereo reconstruction under dynamic traffic conditions: Incorporating height-adaptive vehicle detection and multi-view occlusion optimization," *Automation in Construction*, 2022.

[87] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," in *CVF CVPR*, 2016.

[88] Y. Yao, Z. Luo, S. Li, T. Fang, and L. Quan, "Mvsnet: Depth inference for unstructured multi-view stereo," in *ECCV*, 2018.

[89] F. Wang, S. Galliani, C. Vogel, P. Speciale, and M. Pollefeys, "Patchmatchnet: Learned multi-view patchmatch stereo," in *CVF CVPR*, 2021.

[90] Y. Li and J. Ibanez-Guzman, "Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems," *IEEE Signal Processing Magazine*, 2020.

[91] Y. Chen, J. Zhang, Z. Xie, W. Li, F. Zhang, J. Lu, and L. Zhang, "S-nerf++: Autonomous driving simulation via neural reconstruction and generation," *IEEE TPAMI*, 2025.

[92] T. Qin, C. Li, H. Ye, S. Wan, M. Li, H. Liu, and M. Yang, "Crowd-sourced nerf: Collecting data from production vehicles for 3d street view reconstruction," *IEEE TITS*, 2024.

[93] T. Yuan, Y. Mao, J. Yang, Y. Liu, Y. Wang, and H. Zhao, "Presight: Enhancing autonomous vehicle perception with city-scale nerf priors," in *ECCV*, 2024.

[94] J. Cao, Z. Li, N. Wang, and C. Ma, "Lightning nerf: Efficient hybrid scene representation for autonomous driving," in *IEEE ICRA*, 2024.

[95] L. Caltagirone, M. Bellone, L. Svensson, and M. Wahde, "Lidar–camera fusion for road detection using fully convolutional neural networks," *Robotics and Autonomous Systems*, 2019.

[96] X. Bai, Z. Hu, X. Zhu, Q. Huang, Y. Chen, H. Fu, and C.-L. Tai, "Transfusion: Robust lidar-camera fusion for 3d object detection with transformers," in *CVF CVPR*, 2022.

[97] T. Liang, H. Xie, K. Yu, Z. Xia, Z. Lin, Y. Wang, T. Tang, B. Wang, and Z. Tang, "Bevfusion: A simple and robust lidar-camera fusion framework," *Neurips*, 2022.

[98] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: part i," *IEEE robotics & automation magazine*, 2006.

[99] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "Orb-slam3: An accurate open-source library for visual, visual–inertial, and multimap slam," *IEEE TRO*, 2021.

[100] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping," in *IEEE/RSJ IROS*, 2020.

[101] Q. Li, Y. Wang, Y. Wang, and H. Zhao, "Hdmapnet: An online hd map construction and evaluation framework," in *IEEE ICRA*, 2022.

[102] Y. Liu, T. Yuan, Y. Wang, Y. Wang, and H. Zhao, "Vectormapnet: End-to-end vectorized hd map learning," in *ICML*, 2023.

[103] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social gan: Socially acceptable trajectories with generative adversarial networks," in *CVF CVPR*, 2018.

[104] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data," in *ECCV*, 2020.

[105] J. Ngiam, B. Caine, V. Vasudevan, Z. Zhang, H.-T. L. Chiang, J. Ling, R. Roelofs, A. Bewley, C. Liu, A. Venugopal *et al.*, "Scene transformer: A unified architecture for predicting multiple agent trajectories," *arXiv:2106.08417*, 2021.

[106] R. T. Hay, "Sumo: a history of modification," *Molecular cell*, 2005.

[107] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," in *CoRL*, 2017.

[108] G. Rong, B. H. Shin, H. Tabatabaee, Q. Lu, S. Lemke, M. Možeiko, E. Boise, G. Uhm, M. Gerow, S. Mehta *et al.*, "Lgsvl simulator: A high fidelity simulator for autonomous driving," in *IEEE ITSC*, 2020.

[109] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "Airsim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and service robotics: Results of the 11th international conference*, 2017.

[110] D. Chen, D. Wang, Y. Zhu, and Z. Han, "Digital twin for federated analytics using a bayesian approach," *IEEE IoT-J*, 2021.

[111] Y. Zhang, S. Wang, B. Chen, J. Cao, and Z. Huang, "Trafficgan: Network-scale deep traffic prediction with generative adversarial nets," *IEEE TITS*, 2019.

[112] M. Jiang, Y. Bai, A. Cornman, C. Davis, X. Huang, H. Jeon, S. Kulshrestha, J. Lambert, S. Li, X. Zhou *et al.*, "Scenediffuser: Efficient and controllable driving simulation initialization and rollout," *Neurips*, 2024.

[113] D. S. Lee, A. Karthikeyan, Y. V. Pant, and S. Fischmeister, "Path diffuser: Diffusion model for data-driven traffic simulator," *arXiv:2509.24995*, 2025.

[114] G. Boquet, A. Morell, J. Serrano, and J. L. Vicario, "A variational autoencoder solution for road traffic forecasting systems: Missing data imputation, dimension reduction, model selection and anomaly detection," *Transportation Research Part C*, 2020.

[115] G. Boquet, J. L. Vicario, A. Morell, and J. Serrano, "Missing data in traffic estimation: A variational autoencoder imputation method," in *IEEE ICASSP*, 2019.

[116] X. Pan, D. Seita, Y. Gao, and J. Canny, "Risk averse robust adversarial reinforcement learning," in *IEEE ICRA*, 2019.

[117] Z. Wei, H. Huang, G. Zhang, R. Zhou, X. Luo, S. Li, and H. Zhou, "Interactive critical scenario generation for autonomous vehicles testing based on in-depth crash data using reinforcement learning," *IEEE TIV*, 2024.

[118] R. Zhou, Z. Wei, J. Bian, Q. Yu, S. Tian, H. Huang, G. Gui, and L. Xing, "Marl-based high-risk multi-vehicle scenario generation for autonomous vehicle safety testing," *IEEE IoT-J*, 2025.

[119] W. Ding, H. Lin, B. Li, and D. Zhao, "Causalaf: Causal autoregressive flow for safety-critical driving scenario generation," in *CoRL*, 2023.

[120] Z. Jiang, J. Liu, P. Sun, M. Sang, H. Li, and Y. Pan, "Generation of risky scenarios for testing automated driving visual perception based on causal analysis," *IEEE TITS*, 2024.

[121] D. J. Yeong, G. Velasco-Hernandez, J. Barry, and J. Walsh, "Sensor and sensor fusion technology in autonomous vehicles: A review," *Sensors*, 2021.

[122] D. F. Llorca, R. Hamon, H. Junklewitz, K. Grosse, L. Kunze, P. Seiniger, R. Swaim, N. Reed, A. Alahi, E. Gómez *et al.*, "Testing autonomous vehicles and ai: perspectives and challenges from cybersecurity, transparency, robustness and fairness," *European Transport Research Review*, 2025.

[123] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *CVF CVPR*, 2020.

[124] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine *et al.*, "Scalability in perception for autonomous driving: Waymo open dataset," in *CVF CVPR*, 2020.

[125] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *IJRR*, 2013.

[126] M. Zhou, J. Luo, J. Villella, Y. Yang, D. Rusu, J. Miao, W. Zhang, M. Alban, I. Fadakar, Z. Chen *et al.*, "Smarts: Scalable multi-agent reinforcement learning training school for autonomous driving," *arXiv:2010.09776*, 2020.

[127] E. Vinitsky, N. Lichtlé, X. Yang, B. Amos, and J. Foerster, "Nocturne: a scalable driving benchmark for bringing multi-agent learning one step closer to the real world," *Neurips*, 2022.

[128] S. Ashayer, F. Elhag, P. Sun, and M. Sartipi, "Co-simulation in intelligent transportation systems: A survey targeting connected and autonomous vehicles," *Transportation Research Record*, 2025.

[129] C. Ding and I. W.-H. Ho, "Digital-twin-enabled federated learning and cnn-based channel estimation for urban vehicular channels," *IEEE IoT-J*, 2025.

[130] W. Sun, P. Wang, N. Xu, G. Wang, and Y. Zhang, "Dynamic digital twin and distributed incentives for resource allocation in aerial-assisted internet of vehicles," *IEEE IoT-J*, 2021.

[131] J. Lai, Z. Chen, J. Zhu, W. Ma, L. Gan, S. Xie, and G. Li, "Deep learning based traffic prediction method for digital twin network," *Cognitive Computation*, 2023.

[132] A. Isah, H. Shin, I. Aliyu, R. M. Sulaiman, and J. Kim, "Graph neural network for digital twin network: A conceptual framework," in *ICAIIC*. IEEE, 2024.

[133] H. Zhang, X. Ma, X. Liu, L. Li, and K. Sun, "Gnn-based power allocation and user association in digital twin network for the terahertz band," *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 10, pp. 3111–3121, 2023.

[134] U. Demir, S. Pradhan, R. Kumahia, D. Roy, K. Chowdhury *et al.*, "Digital twins for maintaining qos in programmable vehicular networks," *IEEE Network*, 2023.

[135] R. Pegurri, D. Gasco, F. Linsalata, M. Rapelli, E. Moro, F. Raviglione, and C. Casetti, "Van3twin: the multi-technology v2x digital twin with ray-tracing in the loop," *arXiv:2505.14184*, 2025.

[136] P. Testolina, M. Polese, P. Johari, and T. Melodia, "Boston twin: the boston digital twin for ray-tracing in 6g networks," in *ACM Multimedia Systems*, 2024.

[137] G. Twardokus and H. Rahbari, "Dt-covess: Advancing nextg c-v2x security evaluation through high-fidelity digital twinning," in *IEEE INFOCOM workshop*. IEEE, 2025.

[138] P. Singh, B. Hazarika, K. Singh, W.-J. Huang, and T. Q. Duong, "Genai-enhanced federated multi-agent drl for digital twin-assisted iov networks," *IEEE IoT-J*, 2025.

[139] C. Tan, P. Yu, Z. Qu, L. Zhang, W. Li, X. Qiu, and S. Guo, "Energy-efficient federated learning training optimization for digital twin driven 6g air-ground integrated vehicular networks," *IEEE TITS*, 2025.

[140] A. Das, T. Castiglia, S. Wang, and S. Patterson, "Cross-silo federated learning for multi-tier networks with vertical and horizontal data partitioning," *ACM TIST*, 2022.

[141] L. Zhao, Z. Bi, A. Hawbani, K. Yu, Y. Zhang, and M. Guizani, "Elite: An intelligent digital twin-based hierarchical routing scheme for softwarized vehicular networks," *IEEE TMC*, 2022.

[142] J. Liu, Z. Peng, D. Xu, and Y. Liu, "Revolutionizing wireless modeling and simulation with network-oriented llms," in *IEEE IPCCC*, 2024.

[143] Y. Xie, Q. Wu, P. Fan, N. Cheng, W. Chen, J. Wang, and K. B. Letaief, "Resource allocation for twin maintenance and task processing in vehicular edge computing network," *IEEE IoT-J*, 2025.

[144] L. Tang, A. Wang, B. Xia, Y. Tang, and Q. Chen, "Research on integrated sensing, communication resource allocation and digital twin placement based on digital twin in iov," *IEEE IoT-J*, 2025.

[145] C. Tan, X. Li, L. Gao, T. H. Luan, Y. Qu, Y. Xiang, and R. Lu, "Digital twin enabled remote data sharing for internet of vehicles: System and incentive design," *IEEE TVT*, 2023.

[146] J. Zheng, T. H. Luan, G. Li, Z. Yin, Y. Wu, and M. Dong, "Acdv: Adaptive content delivery for vehicular digital twin networks," *IEEE TVT*, 2025.

[147] K. Zhang, J. Cao, S. Maharjan, and Y. Zhang, "Digital twin empowered content caching in social-aware vehicular edge networks," *IEEE TCSS*, 2021.

[148] Y. Gong, Y. Wei, Z. Feng, F. R. Yu, and Y. Zhang, "Resource allocation for integrated sensing and communication in digital twin enabled internet of vehicles," *IEEE TVT*, 2022.

[149] L. Li, L. Tang, Y. Wang, T. Liu, and Q. Chen, "Intelligent reflecting surface and network slicing assisted vehicle digital twin update," *IEEE TITS*, 2024.

[150] X. Tan, M. Wang, T. Wang, Q. Zheng, J. Wu, and J. Yang, "Adaptive task scheduling in digital twin empowered cloud-native vehicular networks," *IEEE TVT*, 2024.

[151] Z. Zhang, M. Fang, M. Chen, and Y. Liu, "On transferring, merging, and splitting task-oriented network digital twins," in *IEEE MSWiM*. IEEE, 2025.

[152] J. Zheng, T. H. Luan, Y. Hui, Z. Yin, N. Cheng, L. Gao, and L. X. Cai, "Digital twin empowered heterogeneous network selection in vehicular networks with knowledge transfer," *IEEE TVT*, 2022.

[153] J. Zheng, T. H. Luan, Y. Zhang, R. Li, Y. Hui, L. Gao, and M. Dong, "Data synchronization in vehicular digital twin network: A game theoretic approach," *IEEE TWC*, 2023.

[154] X. Chao, J. Jiang, W. Ma, Y. Li, J. Nie, S. Ercisli, and G. Muhammad, "Personalized federated transformer architecture with digital twin for enhanced environmental perception in intelligent iov systems," *IEEE IoT-J*, 2025.

[155] S. S. Ulhas, S. Kannapiran, and S. Berman, "Gan-based domain adaptation for creating digital twins of small-scale driving testbeds: Opportunities and challenges," in *IEEE IV*. IEEE, 2024.

[156] K. Qu and W. Zhuang, "Digital twin assisted intelligent network management for vehicular applications," *IEEE Wireless Communications*, 2024.

[157] Y. Tao, J. Wu, Q. Pan, A. K. Bashir, and M. Omar, "O-ran-based digital twin function virtualization for sustainable iov service response: An asynchronous hierarchical reinforcement learning approach," *IEEE TGCN*, 2024.

[158] A. Paul, K. Singh, C.-P. Li, O. A. Dobre, and T. Q. Duong, "Digital twin-aided vehicular edge network: A large-scale model optimization by quantum-drl," *IEEE TVT*, 2024.

[159] K. Noroozi, T. D. Todd, D. Zhao, and G. Karakostas, "Age of information in digital twin migration," *IEEE TVT*, 2025.

[160] F. Mou, J. Lou, Z. Tang, Y. Wu, W. Jia, Y. Zhang, and W. Zhao, "Adaptive digital twin migration in vehicular edge computing and networks," *IEEE TVT*, 2024.

[161] B. Lu, X. Huang, Y. Wu, L. Qian, D. Niyato, and C. Xu, "Cooperative perception aided digital twin model update and migration in mixed vehicular networks," *IEEE TITS*, 2024.

[162] W. Ding, Z. Yang, M. Chen, Y. Liu, and M. Shikh-Bahaei, "Joint vehicle connection and beamforming optimziation in digital twin assisted integrated sensing and communication vehicular networks," *IEEE IoT-J*, 2024.

[163] K. Zhang, J. Cao, and Y. Zhang, "Adaptive digital twin and multi-agent deep reinforcement learning for vehicular edge computing and networks," *IEEE TII*, 2021.

[164] L. Lin, W. Chen, Q. He, J. Xiong, J. Lin, and L. Lin, "Veco: A digital twin-empowered framework for efficient vehicular edge caching and computation offloading," *IEEE TITS*, 2025.

[165] X. Zhao, J. Zhang, C. Wang, P. Yuan, J. Zhang, and X. Li, "Edt_mtos: An edge digital twin enabled vec multi-hop collaborative task offloading scheme," *IEEE IoT-J*, 2025.

[166] S. Li, S. Li, Y. Sun, B. Wang, B. Wang, and B. Zhang, "Digital twin-assisted computation offloading and resource allocation for multi-device collaborative tasks in industrial internet of things," *IEEE TNSE*, 2025.

[167] H.-N. Yang, J.-Y. Wang, Q. Zeng, M. Cheng, M. Lin, and J.-B. Wang, "Joint task offloading and resource allocation for integrated vlc and sensing in digital twin-aided vehicular edge computing networks," *IEEE TVT*, 2024.

[168] M. Elloumi, G. Kaddoum, and M. Z. Hassan, "Digital twin-empowered interference management for multi-hop internet-of-vehicles networks over millimeter wave bands," *IEEE IoT-J*, 2025.

[169] X. Yuan, J. Chen, N. Zhang, J. Ni, F. R. Yu, and V. C. Leung, "Digital twin-driven vehicular task offloading and irs configuration in the internet of vehicles," *IEEE TITS*, 2022.

[170] L. Tang, L. Wang, D. Fang, and Q. Chen, "Dynamic slice resource management and information synchronization strategy in iov based on digital twin," *IEEE TITS*, 2025.

[171] M. Z. Alam, S. Rahman, M. A. B. Khaled, A. Islam, and A. Jamalipour, "A graph-assisted digital-twin-driven multi-agent shared offloading for internet of vehicles (iovs)," *IEEE IoT-J*, 2025.

[172] J. Yang, F. Lin, C. Chakraborty, K. Yu, Z. Guo, A.-T. Nguyen, and J. J. Rodrigues, "A parallel intelligence-driven resource scheduling scheme for digital twins-based intelligent vehicular systems," *IEEE TIV*, 2023.

[173] L. Zhao, Z. Zhao, E. Zhang, A. Hawbani, A. Y. Al-Dubai, Z. Tan, and A. Hussain, "A digital twin-assisted intelligent partial offloading approach for vehicular edge computing," *IEEE JSAC*, 2023.

[174] X. Zhang, C. Lai, G. Li, and D. Zheng, "Two-phase authentication for secure vehicular digital twin communications," *Computer Networks*, 2025.

[175] K. Wang, J. Dong, S. Wang, Z. Yuan, L. Sha, and F. Xiao, "Rsaka-vdt: Designing reliable and provably secure authenticated key agreement scheme for vehicular digital twin networks," *IEEE TVT*, 2025.

[176] D. Gautam, G. Thakur, P. Kumar, A. K. Das, and Y. Park, "Blockchain assisted intra-twin and inter-twin authentication scheme for vehicular digital twin system," *IEEE TITS*, 2024.

[177] H. Feng, D. Chen, and Z. Lv, "Blockchain in digital twins-based vehicle management in vanets," *IEEE TITS*, 2022.

[178] X. Luo, J. Wen, J. Kang, J. Nie, Z. Xiong, Y. Zhang, Z. Yang, and S. Xie, "Privacy attacks and defenses for digital twin migrations in vehicular metaverses," *IEEE Network*, 2023.

[179] S. Jabeen Siddiqi, A. H. Alobaidi, M. Ahmad Jan, and M. Tariq, "Securing vehicle-to-digital twin communications in the internet of vehicles," *ACM ToMM*, 2025.

[180] C. Oham, R. A. Michelin, R. Jurdak, S. S. Kanhere, and S. Jha, "B-ferl: Blockchain based framework for securing smart vehicles," *Information Processing & Management*, 2021.

[181] R. C. Shit and S. Subudhi, "Ai-powered anomaly detection with blockchain for real-time security and reliability in autonomous vehicles," in *IEEE SPACE*, 2025.

[182] C. Wang, Y. Ming, H. Liu, J. Feng, and N. Zhang, "Secure and flexible data sharing with dual privacy protection in vehicular digital twin networks," *IEEE TITS*, 2024.

[183] M. Ali, G. Kaddoum, W.-T. Li, C. Yuen, M. Tariq, and H. V. Poor, "A smart digital twin enabled security framework for vehicle-to-grid cyber-physical systems," *IEEE TIFS*, 2023.

[184] C. Xu, P. Zhang, X. Xia, L. Kong, P. Zeng, and H. Yu, "Digital twin-assisted intelligent secure task offloading and caching in blockchain-based vehicular edge computing networks," *IEEE IoT-J*, 2024.

[185] X. Wang, M. Wu, M. Hao, D. Ye, J. Kang, and R. Yu, "The butterfly effect in vehicular digital twin systems: Complexity and risk analysis for mixed-traffic scenarios," *IEEE TVT*, 2025.

[186] S. J. Siddiqi, S. Saleh, M. A. Jan, and M. Tariq, "Securing the vetaverse: Web 3.0 for decentralized digital twin-enhanced vehicle–road safety," *Future Generation Computer Systems*, 2025.

[187] H. Feng, H. Lv, and Z. Lv, "Resilience towarded digital twins to improve the adaptability of transportation systems," *Transportation Research Part A*, 2023.

[188] S. Almeaibed, S. Al-Rubaye, A. Tsourdos, and N. P. Avdelidis, "Digital twin analysis to promote safety and security in autonomous vehicles," *IEEE Communications Standards Magazine*, 2021.

[189] H. Zhang, X. Yue, K. Tian, S. Li, K. Wu, Z. Li, D. Lord, and Y. Zhou, "Virtual roads, smarter safety: A digital twin framework for mixed autonomous traffic safety analysis," *arXiv:2504.17968*, 2025.

[190] W. Daluwatta, S. Edirimannage, C. Elvitigala, I. Khalil, and M. Atiquzzaman, "Dt-fu: Digital twin-driven federated unlearning for resilient vehicular networks in the 6g era," *IEEE Communications Magazine*, 2024.

[191] Y. Yigit, L. Maglaras, W. J. Buchanan, B. Canberk, H. Shin, and T. Q. Duong, "Ai-enhanced digital twin framework for cyber-resilient 6g internet-of-vehicles networks," *IEEE IoT-J*, 2024.

[192] M. Palmieri, C. Quadri, A. Fagiolini, and C. Bernardeschi, "Co-simulated digital twin on the network edge: A vehicle platoon," *Computer Communications*, 2023.

[193] L. Tang, Z. Pu, Z. Cheng, D. Fang, and Q. Chen, "Variable granularity vehicle digital twin construction scheme for dt-assisted iovs," *IEEE IoT-J*, 2024.

[194] Z. Wang, X. Liao, X. Zhao, K. Han, P. Tiwari, M. J. Barth, and G. Wu, "A digital twin paradigm: Vehicle-to-cloud based advanced driver assistance systems," in *IEEE VTC*, 2020.

[195] X. Liao, Z. Wang, X. Zhao, K. Han, P. Tiwari, M. J. Barth, and G. Wu, "Cooperative ramp merging design and field implementation: A digital twin approach based on vehicle-to-cloud communication," *IEEE TITS*, 2021.

[196] K. Wu, P. Li, Y. Cheng, S. T. Parker, B. Ran, D. A. Noyce, and X. Ye, "A digital twin framework for physical-virtual integration in v2x-enabled connected vehicle corridors," *IEEE TITS*, 2025.

[197] Z. Wang, O. Zheng, L. Li, M. Abdel-Aty, C. Cruz-Neira, and Z. Islam, "Towards next generation of pedestrian and connected vehicle in-the-loop research: A digital twin co-simulation framework," *IEEE TIV*, 2023.

[198] J. Dong, Q. Xu, J. Wang, C. Yang, M. Cai, C. Chen, Y. Liu, J. Wang, and K. Li, "Mixed cloud control testbed: Validating vehicle-road-cloud integration via mixed digital twin," *IEEE TIV*, 2023.

[199] T. T. M. Tran, C. Parker, and M. Tomitsch, "A review of virtual reality studies on autonomous vehicle–pedestrian interaction," *IEEE THMS*, 2021.

[200] M. Funk Drechsler, J. B. Peintner, G. Seifert, W. Huber, and A. Riener, "Mixed reality environment for testing automated vehicle and pedestrian interaction," in *ACM AutoUI*, 2021.

[201] S. M. Serrano, R. Izquierdo, I. G. Daza, M. A. Sotelo, and D. F. Llorca, "Digital twin in virtual reality for human-vehicle interactions in the context of autonomous driving," in *IEEE ITSC*, 2023.

[202] L. Tang, Z. Cheng, J. Dai, H. Zhang, and Q. Chen, "Joint optimization of vehicular sensing and vehicle digital twins deployment for dt-assisted iovs," *IEEE TVT*, 2024.

[203] E. Karabulut, S. F. Pileggi, P. Groth, and V. Degeler, "Ontologies in digital twins: A systematic literature review," *Future Generation Computer Systems*, 2024.

[204] F. G. Listl, D. Dittler, G. Hildebrandt, V. Stegmaier, N. Jazdi, and M. Weyrich, "Knowledge graphs in the digital twin: A systematic literature review about the combination of semantic technologies and simulation in industrial automation," *IEEE Access*, 2024.

[205] M. Granitzer, V. Sabol, K. W. Onn, D. Lukose, and K. Tochtermann, "Ontology alignment—a survey with focus on visually supported semi-automatic techniques," *Future Internet*, 2010.

[206] M. Dalibor, N. Jansen, B. Rumpe, D. Schmalzing, L. Wachtmeister, M. Wimmer, and A. Wortmann, "A cross-domain systematic mapping study on software engineering for digital twins," *Journal of Systems and Software*, 2022.

[207] L. Lei, K. Zheng, J. Mei, and X. Shen, "Voi-driven joint optimization of control and communication in vehicular digital twin network," *IEEE Network*, 2024.

[208] D. Li, Y. Yang, Y.-Z. Song, and T. Hospedales, "Learning to generalize: Meta-learning for domain generalization," in *AAAI*, 2018.

[209] L. Wang, X. Zhang, H. Su, and J. Zhu, "A comprehensive survey of continual learning: Theory, method and application," *IEEE TPAMI*, 2024.

[210] K. Wang, Z. Li, K. Nonomura, T. Yu, K. Sakaguchi, O. Hashash, and W. Saad, "Smart mobility digital twin based automated vehicle navigation system: A proof of concept," *IEEE TIV*, 2024.

[211] B. Hazarika, K. Singh, C.-P. Li, A. Schmeink, and K. F. Tsang, "Radit: Resource allocation in digital twin-driven uav-aided internet of vehicle networks," *IEEE JSAC*, 2023.

[212] B. Zhang and H. Soh, "Extract, define, canonicalize: An llm-based framework for knowledge graph construction," in *EMNLP*, 2024.

[213] P. Lu, B. Peng, H. Cheng, M. Galley, K.-W. Chang, Y. N. Wu, S.-C. Zhu, and J. Gao, "Chameleon: Plug-and-play compositional reasoning with large language models," *Neurips*, 2023.

[214] Á. Bárkányi, T. Chován, S. Németh, and J. Abonyi, "Modelling for digital twins—potential role of surrogate models," *Processes*, 2021.

[215] S. A. de Parga, J. R. Bravo, N. Sibuet, J. Hernandez, R. Rossi, S. Boschert, E. S. Quintana-Ortí, A. E. Tomás, C. C. Tatu, F. Vázquez-Novoa *et al.*, "Parallel reduced-order modeling for digital twins using high-performance computing workflows," *Computers & Structures*, 2025.

[216] A. Balu, S. Sarkar, B. Ganapathysubramanian, and A. Krishnamurthy, "Physics-aware machine learning surrogates for real-time manufacturing digital twin," *Manufacturing Letters*, 2022.

[217] A. Kontaxoglou, S. Tsutsumi, S. Khan, and S. Nakasuka, "Towards a digital twin enabled multifidelity framework for small satellites," in *PHM Society European Conference*, 2021.

[218] T. Sahin, D. Wolff, M. von Danwitz, and A. Popp, "Towards a hybrid digital twin: physics-informed neural networks as surrogate model of a reinforced concrete beam," *arXiv*, 2024.

[219] M. G. Fernández-Godino, "Review of multi-fidelity models," *arXiv*, 2016.

[220] D. Chen, E. Yurtsever, K. A. Redmill, and Ü. Özgüner, "Using collision momentum in deep reinforcement learning based adversarial pedestrian modeling," in *IEEE IV*, 2023.

[221] W. Ding, C. Xu, M. Arief, H. Lin, B. Li, and D. Zhao, "A survey on safety-critical driving scenario generation—a methodological perspective," *IEEE TITS*, 2023.

[222] S. Feng, H. Sun, X. Yan, H. Zhu, Z. Zou, S. Shen, and H. X. Liu, "Dense reinforcement learning for safety validation of autonomous vehicles," *Nature*, 2023.

[223] Y. Wang, H. Huang, Z. Xu, K. Sun, B. Guo, and J. Zhao, "Learning from risk: Llm-guided generation of safety-critical scenarios with prior knowledge," *arXiv preprint arXiv:2511.20726*, 2025.

[224] F. Kargl, A. Friedman, and R. Boreli, "Differential privacy in intelligent transportation systems," in *ACM WiSec*, 2013.

[225] N. Bertozzi, A. Geraci, L. Bergamasco, E. Ferrera, E. Pristeri, C. Pastrone *et al.*, "A distributed event-orchestrated digital twin architecture for optimizing energy-intensive industries." in *IoTBDS*, 2025.

[226] M. Nilsson, O. Schelén, A. Lindgren, U. Bodin, C. Paniagua, J. Delsing, and F. Sandin, "Integration of neuromorphic ai in event-driven distributed digitized systems: Concepts and research directions," *Frontiers in Neuroscience*, 2023.

[227] M. Neu, J. Becker, P. Dorwarth, T. Ferber, L. Reuter, S. Stefkova, and K. Unger, "Real-time graph building on fpgas for machine learning trigger applications in particle physics," *Computing and Software for Big Science*, 2024.

[228] S. M. Raza, R. Minerva, N. Crespi, M. Alvi, M. Herath, and H. Dutta, "A comprehensive survey of network digital twin architecture, capabilities, challenges, and requirements for edge–cloud continuum," *Computer Communications*, 2025.

[229] 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.

[230] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, "Fourier neural operator for parametric partial differential equations," *arXiv preprint arXiv:2010.08895*, 2020.

[231] Y.-Z. Lei, Y. Gong, D. Chen, Y. Cheng, and X. T. Yang, "Reconstructing physics-informed machine learning for traffic flow modeling: a multi-gradient descent and pareto learning approach," *Transportation Research Part C: Emerging Technologies*, 2025.

[232] Lei, Yuan-Zheng and Gong, Yaobang and Chen, Dianwei and Cheng, Yao and Yang, Xianfeng Terry, "Potential failures of physics-informed machine learning in traffic flow modeling: theoretical and experimental analysis," *arXiv preprint arXiv:2505.11491*, 2025.