International Conference on Science, Engineering & Management Trends

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

Digital Twin in the 6G Internet of Vehicles: A Concise Review of Channel Modelling, Learning, and Security

Rashmi Vanajakar

Research Scholar, Srinivas Institute of Engineering and Technology, Srinivas University, Mukka, Mangalore, Karnataka, India and Lecturer Selection Grade-1, Department of Electronics and Communication, 121-Government Polytechnic for Women, Hubli.

Abstract- The Internet of Vehicles is being recast around two technologies that are reaching deployment maturity in the same window. On the cellular side, sixth-generation networks open terahertz spectrum, integrated sensing and communication, and a space-air-ground integrated fabric in which roadside units, unmanned aerial vehicles, and low-earth-orbit satellites all serve as edge servers. Digital twin networks add a continuously synchronised virtual counterpart for every vehicle, road segment, and radio environment, on which learning algorithms can operate as if it were the physical network itself. Each technology has a substantial literature on its own; the joint deployment they are now becoming raises four questions that no single paper resolves: how a radio-frequency twin is grounded in the physics of a millimetre-wave channel, how vehicle twins migrate across heterogeneous edge servers, how learning is split between vehicles and their twins without leaking data, and how the resulting pipeline is defended against active falsification. This review reads fifteen recent peer-reviewed frameworks against those questions and treats them as components of a single deployed pipeline rather than as isolated proposals. The 3D ray-tracing RF digital twin scheme of Liu et al. is given a dedicated discussion because most of the learning-oriented works reviewed here rest on a channel abstraction that, in deployment, would have to be served by some form of RF twin. Two comparison tables consolidate the readings, and five research gaps are identified: unverifiable twin fidelity, fragmented benchmark practice, opaque synchronisation cost, weak active-adversary threat models, and absent end-to-end energy accounting. Each gap is paired with an incremental, testable next step.

Keywords- 6G networks, Internet of Vehicles, digital twin, vehicle twin, millimetre-wave channel modelling, 3D ray tracing, reconfigurable intelligent surface, federated learning, multi-agent reinforcement learning, ISAC, SAGIN, edge intelligence, V2DT security

I. INTRODUCTION

Two technological tracks have arrived at the Internet of Vehicles inside the same standardisation cycle. Sixth-generation cellular standardisation is advancing through ITU-R IMT-2030 [1] and 3GPP Release

DOI: http://doi.org/10.5281/zenodo.20557590

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

20 [2], with early prototypes already exercising spectrum above 100 GHz, integrated sensing and communication, reconfigurable intelligent surfaces, and a space-air-ground integrated fabric in which terrestrial base stations are joined by unmanned aerial vehicles and low-earth-orbit satellites acting as edge servers [3]. Digital twin networks have moved in parallel from manufacturing-era visualisation toward a working definition that includes lifecycle synchronisation, federated model exchange, and safe actuation back on the live network [4]. The Internet of Vehicles (IoV) is the vertical in which the two technologies now collide most visibly, and a 6G-enabled IoV is expected to support autonomous driving, cooperative perception, and high-definition video analytics under millisecond-scale latency [5].

Vehicular networks raise the cost of getting any part of this wrong. A vehicle whose twin has drifted from its physical state can take an offloading decision that wastes a radio resource block; a vehicle whose twin has been tampered with can take a decision that puts the occupant at risk. The same property that makes a digital twin attractive, namely that it concentrates state from many sensors into a single addressable replica, also makes the twin a high-value target for adversaries and a high-consequence dependency for the safety case.

Recent work has responded along five complementary axes. A first body of work has concentrated on the radio-frequency twin itself, asking how the channel that connects a vehicle to its twin should be modelled when the carrier sits in the millimetre-wave band, when scatterers are moving objects rather than static buildings, and when an aerial reconfigurable intelligent surface is used to shape the propagation. The work of Liu et al. [6] is the most direct expression of this direction and is discussed in detail in Section 2.2. A second body has concentrated on the migration and synchronisation of vehicle twins across heterogeneous edge servers, recognising that the twin must follow the vehicle [7], [8]. A third body has placed the digital twin between the vehicle and a learning system, using federated learning, transformer pre-training, deep reinforcement learning, generative AI, and most recently quantum DDPG to drive offloading, caching, slicing, and beamforming decisions on the twin rather than on the physical vehicle [9]-[10]. A fourth body has begun the work of producing reference architectures and definitions that the others can be evaluated against [11], [12]. A fifth body has taken the step the empirical papers usually do not, and asked how the vehicle-to-twin channel itself can be secured against false data injection [13].

Read paper by paper, the picture is one of progress; read across all five strands, the picture is less consistent than any individual paper suggests. Channel-level work validates against finite-difference time-domain ground truth but assumes that the twin is available at the resolution the learning frameworks would need. Learning frameworks report convergence and energy savings under simulated channels whose fidelity is left implicit. Architecture papers describe the lifecycle of a network twin without connecting it to the physical-layer measurements that would populate it. The security paper reports tamper-detection rates close to 100 percent but is evaluated against a small family of synthetic attacks. The empirical and architectural strands do not yet share a benchmark.

Motivation and Contribution

A reader trying to plan a digital twin component for a 6G-IoV deployment today has to move between sub-literatures that rarely cite each other and whose evaluations are not directly comparable. The dominant pattern is single-scheme, single-dataset reporting against bespoke baselines, in which a learning result is decoupled from the channel it implicitly assumes and the threat model is named only in a closing paragraph. This review is built to be useful to that reader.

The contributions of this review are five.

The reviewed methods are organised along the deployment-pipeline axis: physical-layer and architecture frameworks that fix the substrate on which a digital twin runs (channel models, reference architectures, cross-layer orchestration), and learning-and-resource frameworks that operate on top of

DOI: http://doi.org/10.5281/zenodo.20557590

that substrate (twin migration, slice and resource allocation, federated perception, vehicle-to-twin security).

Frameworks are compared not only on headline metrics such as average delay or detection accuracy but also on the data they were evaluated against, the threat model they actually assume, the cost of maintaining twin synchronisation, the role of edge servers outside the terrestrial network, and the reproducibility of the reported results.
A dedicated subsection is given to the RF digital twin scheme of Liu et al. [6]. Most of the learning-oriented frameworks reviewed here depend implicitly on a channel abstraction that, in deployment, would have to be served by some form of RF twin; the work in [6] is the most concrete recent attempt to construct that abstraction from physics rather than from a pathloss exponent.

Two comparison tables place the surveyed primary works on a common grid: a physical-layer and architecture table, and a learning, resource, and security table. The columns are uniformly defined so that schemes can be compared on the dimensions a deployment engineer actually faces.

The research gap is discussed under five explicit headings, namely unverifiable twin fidelity, fragmented benchmark practice, opaque synchronisation cost, weak active-adversary threat models, and absent end-to-end energy accounting. The future-direction section frames the next steps as incremental and testable rather than as a single integrated grand solution.

The rest of the paper is organised as follows. Section 2 reviews the recent literature in five thematic subsections, with Liu et al. [6] given its own subsection inside the channel sub-area, and consolidates the readings in two comparison tables. Section 3 discusses the research gap and the future directions that follow from it. Section 4 concludes.

II. RELATED WORK

Recent work on digital twin in 6G-enabled IoV can be split into five strands: reference architectures and definitions; channel-level RF twin and physical-layer modelling; twin migration, caching, and synchronisation across edge servers; learning-driven slice, resource, and perception management; and vehicle-to-twin security. The subsections that follow read each strand in turn. The work of [6] is placed inside the channel sub-area because that is the role it plays in the wider pipeline.

2.1 Reference Architectures and Orchestration for the DT-IoV

Reference architectures for the network digital twin (NDT) have lagged the empirical literature, and the two recent contributions that try to close the gap are very different in scope. Zaki-Hindi et al. [11] consolidate fragmented industry and standardisation efforts into a single working definition of the NDT in 6G, then propose a reference functional architecture decomposed into application, physical, digital, and management domains, with explicit interfaces and a harmonised data model. The contribution is organised around lifecycle-aware procedures, namely instantiation, synchronisation, model selection and update, federation, and safe actuation, and ties AI and simulation workflows together through MLOps and co-simulation. The decomposition is useful in its own right because it makes the synchronisation interface a first-class element, which the learning-heavy frameworks reviewed below tend to leave implicit.

Li [12] provides a complementary overview from a communications angle, surveying the key technologies of digital twin networks in the 6G era and proposing an end-to-end twin architecture for non-terrestrial networks. The treatment of channel modelling, AI models, and intelligent operation and maintenance is at the textbook level, and the paper is best read as a structuring guide rather than as a source of new methods. Both works are hesitant on the question of how a NDT is verified against

DOI: http://doi.org/10.5281/zenodo.20557590

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

the live network it claims to mirror, which is the gap most directly exploited by the security paper [13] discussed in Section 2.5.

Cross-layer orchestration takes the reference picture and tries to make it operational. Farre et al. [14] propose a digital twin-driven cross-layer orchestration framework for 6G UAV and non-terrestrial networks that runs three optimisation tiers on separate timescales: a strategic tier driven by a genetic algorithm, a tactical tier driven by proximal policy optimisation and multi-agent reinforcement learning, and a clustering layer that uses K-means for traffic load balancing. The twin itself is exposed as a three-tier physical-digital-application stack. The reported gains on reference signal received power, signal-to-interference-plus-noise ratio, throughput, and physical resource block utilisation are non-trivial, although the validation is confined to simulation and the timescale boundary between the strategic and tactical tiers is set by design rather than learned from workload.

Hu et al. [10] take a service-oriented view in a healthcare IoE setting, driving slice orchestration through reinforcement learning whose state includes large-AI-model semantic features alongside the DT predictions. The reported 42 to 43 percent reduction in service-level-agreement violations against an RL-only baseline is large, but the LAM is treated as a black box. The setting is not vehicular, but the design pattern of twin state plus semantic features as RL input recurs in the IoV frameworks reviewed below.

RF and Channel Modelling for the DT-IoV

Channel modelling has been the most physics-facing strand of the literature, and the work of Liu, Sun, Marine, and Wu [6] is the recent contribution that goes furthest in tying the digital twin to the radio-frequency domain rather than to a statistical pathloss model. Liu et al. propose an RF digital twin in which a three-dimensional ray-tracing model is constructed for the millimetre-wave channel between vehicles, base stations, and an aerial reconfigurable intelligent surface, referred to as an AIRS. The 3D imaging of the vehicle is built up from scattering sites obtained by Fourier transform of the impulse response, so that the twin captures the precise spatial distribution of strong scatterers rather than treating the vehicle as a point reflector. The channel is then re-evaluated for each new geometry through 3D ray tracing, with Doppler shift and micro-Doppler effects included for moving objects.

The mathematical model is set up over a rectangular target area on the x-y plane with a uniform planar AIRS array of size N = Nx by Ny placed at altitude H, and a uniform planar transmit array of size M = My by Mz at the source node. Closed forms are given for the source-to-AIRS and AIRS-to-user power gains under free-space line-of-sight, and for the array response vectors with phase delays psi(nx,ny). The received signal collapses to the familiar y = h^H Theta G v sqrt(P) s + n form, where Theta is the diagonal phase-shift matrix of the AIRS, and the optimisation is formulated as max-min SNR over the AIRS position, the phase configuration, and the source beamforming vector. The simulation is carried out on a 100 GHz carrier with 800 MHz bandwidth, a 512-element cross-polarised array (16 by 16 by 2), EIRP up to 60 dBm, and pathloss exponent three. Four scenarios are compared against FDTD ground truth, and the reported SINR and throughput curves track the FDTD reference closely. The gain from a 256-element MIMO array over a 16-element baseline is substantial at the 100-metre distance.

The strengths of [6] are direct. The twin is grounded in the physics of the channel rather than in a pathloss exponent. The validation against FDTD makes the modelling error visible scenario by scenario rather than as an aggregate number. The inclusion of an aerial RIS folds 6G-era network topology into the same model as the propagation, which is what a deployment of the DT-IoV would actually look like. Energy efficiency is framed as the outer optimisation problem, consistent with the wider green-IoV literature [15] and with task-offloading work that uses an IRS configuration on top of a DT-driven scheduler [16].

DOI: http://doi.org/10.5281/zenodo.20557590

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

The limitations are equally direct. The 3D ray-tracing pipeline is computationally heavy because the correctness of the result depends on the fidelity of the geometric model, and the colour-based FDTD modelling of the vehicle from 3Dmax assets is a manual step that does not yet scale to a city-level twin. The energy-efficient optimisation is implicitly assumed to run on the twin rather than on the vehicle, but the cost of maintaining the twin at the resolution the ray tracer needs is not separately accounted for. The threat model is absent: the channel twin is treated as a faithful mirror of the propagation, with no consideration of an adversary who falsifies the geometry, the scatterer map, or the AIRS phase profile. None of these limitations contradict the contribution; they mark the boundary along which the work in [6] meets the learning-oriented frameworks of Section 2.4 and the security-oriented framework of Section 2.5.

Twin Migration, Caching, and Synchronisation Across Edge Servers

Once the twin has been instantiated, it has to follow the vehicle, and ground roadside units do not provide uniform coverage. Yin et al. [7] address the resulting twin-migration routing problem in a 6G-enabled IoV in which RSUs, UAVs, and satellites all act as edge servers for vehicle-twin construction and update. Three challenges are named: non-uniform deployment of ground edge servers, dynamic workload caused by vehicle mobility, and the calculation complexity of multi-objective migration. The response is two coordinated components. An LSTM-augmented Transformer is used to predict the long-term workload of edge servers, and a diffusion-model multi-agent PPO algorithm is used to obtain near-optimal migration routes that balance latency, energy, and reliability. The reported gains over MAPPO, MADDPG, and Greedy baselines are consistent across the workload-prediction and migration-routing tasks. The contribution is one of the few that explicitly integrates non-terrestrial edge servers into the twin lifecycle rather than treating them as an afterthought.

Zeng et al. [8] address a closely related sub-problem under a different learning paradigm. Their DAPR framework combines digital twin with asynchronous federated learning and deep reinforcement learning for predictive edge caching in vehicular edge computing. The client-selection step uses mobility prediction and data-quality assessment so that vehicles likely to leave RSU coverage or carrying low-quality data do not stall asynchronous aggregation. The predictive model couples a variational autoencoder with a gated recurrent unit, the first to capture latent data distribution and the second to capture temporal structure. Evaluation is on the Beijing taxi GPS set with 10,357 cabs and on MovieLens 1M, and the reported gains in cache-hit ratio, transmission delay, and cumulative reward over federated-learning and reinforcement-learning baselines are non-trivial. The cost is that mobility prediction has to be accurate at the timescale of an RSU dwell, and the framework reports its prediction loss in aggregate rather than in worst-case terms.

Hui et al. [17] sit in the same neighbourhood but address a different decision, namely how a DT-enabled edge collaboration scheme should handle composite services in autonomous vehicular networks. Multiple service types are recognised, and the long-tail service request groups form coalitions with DT service providers through a coalition game; resource pricing and purchase are then negotiated through a Stackelberg game. The Stackelberg-plus-coalition combination is one of the cleaner game-theoretic readings of the edge-collaboration problem in the recent IoV literature, and the explicit separation between long-tail and bursty request groups is realistic.

Learning-Driven Slice, Resource, and Perception Management

The bulk of the recent DT-IoV literature places a learning algorithm on top of the twin and uses the twin as a queryable substrate for offloading, slicing, resource allocation, or perception. The algorithms differ but a common pattern is visible: the twin supplies a state that includes channel quality, queue length, and predicted mobility, the agent acts on offloading or slice configuration, and the reward is a weighted sum of latency, energy, and a quality-of-service penalty.

DOI: http://doi.org/10.5281/zenodo.20557590

Zhan et al. [9] take this pattern into the ISAC-enabled IoV setting. The decision is long-term task offloading and resource allocation, and the objective is the long-term average system cost defined as a weighted combination of delay and energy under queue stability. Lyapunov optimisation is used to decompose the long-horizon stochastic problem into a sequence of per-slot subproblems with bounded backlog, after which a digital-twin-assisted multi-agent PPO algorithm solves the per-slot decisions. The simulation results show faster convergence, better stability, and lower long-term cost against MAPPO, DDPG, and Lyapunov-plus-heuristic baselines. The Lyapunov-plus-PPO design is the methodologically careful part of the paper because it gives an explicit handle on queue stability that an ad-hoc DRL reward cannot. The trade-off is that the V parameter controlling the cost-stability balance is treated as a tunable rather than learned, and the ISAC channel is parameterised through pathloss rather than through a ray-traced twin in the sense of [6].

Tang et al. [18] move the decision up one tier, into network slicing. The problem is joint dual time-domain slice resource management and DT deployment in the IoV. The long time-domain decision is DT deployment, addressed through a Q-value-based deep transfer reinforcement learning algorithm, and the short time-domain decision is per-slot resource allocation plus DT synchronisation weight adjustment, addressed through an LSTM-augmented multi-agent PPO. DT utility is quantified along three dimensions, namely completeness, load contribution, and resource reliability, and a price incentive mechanism is added to balance supply and demand. The dual-timescale design is one of the more careful treatments of the trade-off between DT placement (slow) and slice allocation (fast), and the explicit utility decomposition is useful because most other frameworks fold synchronisation cost into a single weight.

Wang et al. [19] integrate DT and generative AI in a UAV-assisted setting. Their three-layer DTG-HACA architecture combines a DT layer for real-time vehicle/UAV state synchronisation and trajectory simulation, a high-altitude-platform layer for low-latency offloading and solar-powered endurance, and a physical layer for UAV-vehicle-RSU edge collaboration. UAV trajectory optimisation uses MADDPG with prioritised experience replay to balance communication overhead and energy. Joint edge caching and task offloading under privacy constraints is handled by a federated deep reinforcement learning scheme augmented with a generative adversarial network. The reported five-metric evaluation, covering training stability, computational capacity, offloading efficiency, cache-hit rate, and energy consumption, is one of the broader validation grids in the recent literature. The GAN component is used for data augmentation under partial information; whether it would survive an adversarial input is not tested.

Ansere et al. [20] sit further out on the algorithmic axis. Their quantum machine learning DDPG for digital-twin semantic vehicular networks integrates quantum superposition and annealing into the actor-critic loop to widen the effective action search, and pairs the resulting controller with a DT network and an IRS for joint semantic optimisation. The reported gains over classical DDPG on latency and energy in dynamic IoV settings are visible, but the contribution is constrained by the present state of quantum hardware. The paper is best read as a forward-looking position that makes the semantic-decision boundary explicit rather than as a deployable system.
Selvaraj et al. [21] propose the Distributed Intelligence Framework (DIF) for 6G autonomous transport systems, which combines federated learning with DT-on-edge simulation. FL trains a global perception model across autonomous vehicles while DT on edge servers supplies real-time simulation that augments local data. Differential privacy is applied to local updates, secure aggregation defends against model poisoning and eavesdropping, and adversarial update detection filters outliers at the aggregator. The evaluation reports stable convergence across edge servers of different floating-point capacities, which is a more honest test than the usual single-server convergence curve. The framework

DOI: http://doi.org/10.5281/zenodo.20557590

is one of the few that explicitly attempts to bring DT and FL defences into the same pipeline, and the differential-privacy budget is reported, which the bulk of the FL-in-IoV literature still leaves unstated.

Chao et al. [22] approach perception itself rather than resource management. DT-Trans is a personalised federated transformer architecture whose global model, the twin-enhanced vision transformer, is pre-trained on massive synthetic DT data and then fine-tuned through low-rank adaptation on real-world vehicle data. The CD-PFL-Trans algorithm splits the transformer into a shared encoder and client-specific decoder heads, and groups vehicles with similar driving patterns through hierarchical clustering. The design directly addresses the well-known sim-to-real gap of synthetic DT pre-training and is one of the cleanest recent uses of parameter-efficient adaptation in the IoV setting. The trade-off is the cost of constructing the synthetic DT data set at sufficient diversity, which is reported through dataset size but not through diversity metrics.

Security and Privacy of Vehicle-to-Twin Communication

Almost every learning framework discussed above assumes that the data flowing between the vehicle and the twin is faithful. Siddiqi et al. [13] take exactly the opposite starting point. Their study examines the vulnerability of vehicle-to-digital-twin (V2DT) communication to false data injection attacks (FDIAs) and proposes a defence that combines blockchain-based decentralised storage of vehicle dynamics data with a deep-learning sensor-analysis layer. The pipeline stages the vehicle data through a blockchain buffer en route to the edge DT, after which the deep-learning model classifies anomalous driving and FDIA events. The reported figures, 100 percent tamper detection in V2DT communication and 96 percent classification accuracy for anomalous behaviour, are high enough that the more interesting question is how the model generalises beyond the synthetic FDIA classes it is trained on, which the paper does not directly answer. The contribution is the framing more than the figure: it makes the V2DT channel an explicit attack surface that the rest of the literature reviewed here would otherwise leave un-modelled. Selvaraj et al. [21], although primarily a learning framework, is the other paper in the reviewed set that engages with the defensive surface, but the federated training loop it secures covers a different part of the attack space than the V2DT channel that [13] secures.

Cross-Scheme Reading

Tables 1 and 2 place the surveyed primary works on a common grid. The work of [6] is the only contribution that supplies the channel twin from physics; every other scheme borrows the channel from a pathloss model or a simulation oracle. The reference architectures [11], [12] supply the lifecycle vocabulary that the empirical frameworks would otherwise have to invent, and the cross-layer orchestration paper [14] is the most ambitious recent attempt to operationalise that vocabulary, although it does so without a physical channel twin. The twin-migration and caching works [7], [8] address the part of the lifecycle the architecture papers leave abstract, namely how the twin actually follows the vehicle. The learning side shows a converging design pattern but diverging algorithmic choice. Multi-agent reinforcement learning on a DT state with a weighted latency-energy reward is the dominant template [9], [18], [19], [14], and the variation is in how stability is enforced (Lyapunov in [9], explicit utility decomposition in [18], hierarchical decomposition in [14]) and in how the twin is introduced to the agent (state augmentation in [9], [18], synthetic pre-training in [22], game-theoretic exchange in [17]). Security and architecture sit at opposite ends of the maturity axis: the reference architecture papers make the lifecycle explicit but leave the threat model abstract [11], [12], while [13] makes the threat model explicit but exists outside the architectural framing.

Table 1. Physical-layer and reference-architecture frameworks for the DT-IoV.

Ref.	Contribution	Twin grounding / decomposition	Validation	Limitation

DOI: http://doi.org/10.5281/zenodo.20557590

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

[6]	RF DT with 3D ray tracing for mmWave IoV (Liu et al., 2024)	Physics: 3D scattering map, ray tracing, AIRS RIS, 256-element MIMO	FDTD ground truth, 100 GHz, 800 MHz BW, four scenarios	Ray-tracing compute cost; manual geometric modelling; no threat model
[11]	Reference functional architecture for NDT in 6G (Zaki-Hindi et al., 2026)	Application / physical / digital / management; lifecycle procedures	Early implementations of complementary assets	No quantitative twin-to-network fidelity metric
[12]	Survey of DT technology for future 6G (Li, 2025)	End-to-end DT-NTN reference	Conceptual	Textbook-level treatment; no new methods
[14]	DT-driven cross-layer orchestration for 6G UAV-NTN (Farre et al., 2026)	Three tiers: physical / digital / application; strategic + tactical RL	Simulation on RSRP, SINR, throughput, PRB	Hand-set timescale boundary; simulation-only
[10]	LAM-enhanced DT-driven 6G healthcare slicing (Hu et al.)	Physical / digital / application; LAM features as RL state	42-43% SLA-violation reduction vs RL-only baseline	LAM treated as black box; not vehicular

Table 2. Learning, resource-management, and security frameworks for the DT-IoV.

Ref.	Method	Decision target	Edge substrate / dataset	Reported gain
[7]	LSTM-Transformer + DM-MAPPO	Twin migration routing	RSU + UAV + LEO; synthetic 6G-IoV	Lower latency/energy vs MAPPO, MADDPG, Greedy
[8]	DT + asynchronous FL + DRL (DAPR)	Predictive edge caching with mobility-aware client selection	Beijing taxi GPS (10,357 cabs), MovieLens 1M	Higher cache-hit ratio, lower transmission delay
[9]	Lyapunov + DT-MAPPO	Long-term offloading and resource allocation in ISAC IoV	RSU + dual-transmission vehicles; synthetic ISAC	Lower long-term cost vs MAPPO, DDPG, Lyapunov-heuristic
[21]	FL + DT + DP + secure aggregation (DIF)	Cooperative perception under privacy	Heterogeneous edge servers; synthetic ATS	Stable convergence under capacity heterogeneity; explicit DP budget
[22]	TE-ViT + CD-PFL-Trans personalised FL	Environmental perception under FL	Synthetic DT data + real driving data	Accuracy gain via LoRA + clustered personalisation
[18]	QDTRL + LSTM-MAPPO dual timescale	DT deployment (LTD) + slice allocation (STD)	Multi-tenant slice over IoV; synthetic slice mix	Lower QoS violation; DT utility quantified along three axes
[17]	Coalition + Stackelberg games	Composite service edge collaboration	DT-L-SRG + DT-SP + CBS; synthetic AVN	Higher utility vs OPOP/RPOP/OPRP baselines

DOI: http://doi.org/10.5281/zenodo.20557590

[19]	MADDPG-IPER + FDRL-GAN (DTG-HACA)	UAV trajectory + caching + offloading under privacy	RSU + UAV + HAPS; synthetic UAV-IoV	Five-metric improvement
[20]	Quantum DDPG	DT semantic optimisation under IRS	Vehicles + IRS; synthetic IoV	Lower latency, higher EE vs classical DDPG
[13]	Blockchain-buffered V2DT channel + DL anomaly	FDIA detection in V2DT communication	Edge DT + RSU; synthetic FDIA dataset	100% tamper detection, 96% anomaly classification

III. RESEARCH GAP AND FUTURE DIRECTIONS

Reading the two tables side by side surfaces the gaps that any individual paper, on its own, tends to hide. Five are worth naming explicitly.

The first gap is unverifiable twin fidelity. The work of [6] validates its 3D ray-traced channel against FDTD ground truth scenario by scenario, which is the most explicit fidelity check in the reviewed set. Every learning framework above assumes a channel oracle of similar quality but never measures the deviation between its assumed channel and the channel the deployed twin would actually serve. Reference-architecture work [11] names synchronisation and model-selection procedures but does not specify a fidelity metric, and the survey [12] leaves the verification interface implicit. A learning result reported under an unverified oracle is hard to interpret: an energy saving of 20 percent against a baseline that shares the same oracle does not transfer if the oracle itself drifts at deployment.

The second gap is fragmented benchmark practice. Twin migration is evaluated on synthetic 6G-IoV traces [7]; caching is evaluated on the Beijing taxi GPS set and MovieLens 1M [8]; offloading is evaluated on synthetic ISAC channels [9]; slicing is evaluated on a synthetic multi-slice mix [18]; perception is evaluated on synthetic DT data plus a real driving set [22]; FDIA detection is evaluated on a synthetic attack family [13]. No single benchmark exercises the full pipeline at once, so two papers that claim improvements on the same metric on different datasets cannot be ranked against each other today.

The third gap is opaque synchronisation cost. The twin is reported as a state input in [9], [18], [19], [14] but the cost of maintaining the twin at the resolution those agents call is not separately accounted for. The work of [6] is honest about this cost at the channel level; Tang et al. [18] is the only learning framework that explicitly factors synchronisation weight into the optimisation. A learning gain that presupposes a continuously refreshed channel twin will look very different when the twin is refreshed at the rate a budget-constrained edge server can actually afford.

The fourth gap is weak active-adversary threat models in the learning frameworks. Siddiqi et al. [13] is the one paper in the set that designs against an active adversary on the V2DT channel; Selvaraj et al. [21] is the one paper that designs against active adversaries in the federated training loop. The remaining learning frameworks assume an honest physics layer and an honest training loop. A twin in deployment is exactly the kind of high-value target an adversary would attempt to manipulate, and a reward written in terms of a falsifiable state can be made arbitrary by an attacker who can write that state.

The fifth gap is absent end-to-end energy accounting. Energy efficiency is the framing objective of [6] and is reported as a metric in [7] to [9], [18], [19], [20]. Every paper reports the energy of the decision it optimises, but no paper accounts for the energy of running the twin itself, of running the learning

DOI: http://doi.org/10.5281/zenodo.20557590

algorithm, of communicating updates to and from the aggregator, and of validating the twin against the live network. A 20 percent radio-energy saving that hides the cost of running a synthetic transformer pre-training pipeline on a city-scale twin is not a 20 percent saving.

The directions that follow from these gaps are incremental and testable. A shared 6G-IoV evaluation suite that combines the channel level of [6] with the mobility and offloading levels of [7] to [19] and the FDIA surface of [13] would let frameworks be compared on a common grid. A twin-fidelity protocol that reports a worst-case rather than mean deviation between the twin and the live network, in the spirit of the FDTD-against-raytrace comparison of [6], would give a deployment engineer the verification interface that [11] currently leaves abstract. A standard synchronisation-cost line item, expressed as energy and update bandwidth per second per vehicle twin, would close the gap between the radio-level energy accounting of [6] and the system-level energy savings reported in the learning frameworks. An active-adversary evaluation layered on top of the existing FL+DT pipelines, in the spirit of [21] but adopting the V2DT attack family of [13], would expose how robust the convergence and accuracy gains actually are. Beyond these incremental moves, the wider research substrate is framed by the digital-twin-enabled 6G outline of [23], the 6G vision of [24], and the AI-empowered 6G roadmap of [25]. The algorithmic substrate of the learning frameworks reviewed above is itself a mix of federated learning at the McMahan-style aggregator [26], parameter-efficient adaptation through low-rank adapters [27] (a direction [22] already pursues), and the reinforcement-learning primitives DDPG [28] and PPO [29], applied on top of RIS-aided propagation in the tutorial sense of [30]. Vision transformers and self-supervised encoders are likely to displace CNN backbones for perception, and large AI models are entering the orchestration loop as in [10].

IV. CONCLUSION

This review has read the recent 6G-enabled Internet-of-Vehicles digital-twin literature as a single deployed pipeline whose stages, namely channel, lifecycle, twin migration, learning-driven slice and resource control, and vehicle-to-twin security, are usually presented in isolation. The 3D ray-tracing RF digital twin scheme that anchored this review is the most concrete recent attempt to ground the channel stage of that pipeline in physics rather than in a pathloss exponent, and the reference-architecture line of work is the most concrete recent attempt to ground the lifecycle stage in interfaces rather than in slogans. The learning-oriented frameworks reviewed here supply increasingly sophisticated decision policies on top of the twin, but they share an assumption that the twin is faithful, that the cost of maintaining the twin is folded into a single weight, and that the threat model can be left to the closing paragraph. The security paper is the rare contribution that takes that assumption apart. Read across all five strands, the field has made measurable progress at the layer level and has not yet produced an end-to-end benchmark that joins the layers. The next useful steps are the incremental ones named in the preceding section: a shared evaluation suite, a twin-fidelity protocol, a synchronisation-cost line item, an active-adversary evaluation, and an end-to-end energy-accounting line. Taken as a set, the works reviewed here are the components a deployment-oriented agenda would have to combine, and the gaps named above mark where the combination still leaves the engineer without an answer.

REFERENCES

1. ITU-R, "Framework and overall objectives of the future development of IMT for 2030 and beyond (IMT-2030)," Recommendation ITU-R M.2160, Nov. 2023.
2. 3GPP, "Release 20 description," 3GPP TR 21.920, 2025.
3. W. Jiang, B. Han, M. A. Habibi, and H. D. Schotten, "The road towards 6G: A comprehensive survey," IEEE Open J. Commun. Soc., vol. 2, pp. 334-366, 2021.

DOI: http://doi.org/10.5281/zenodo.20557590

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

4. S. Mihai et al., "Digital twins: A survey on enabling technologies, challenges, trends and future prospects," IEEE Commun. Surveys Tuts., vol. 24, no. 4, pp. 2255-2291, 4th Quart. 2022.
5. M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorzi, "Toward 6G networks: Use cases and technologies," IEEE Commun. Mag., vol. 58, no. 3, pp. 55-61, Mar. 2020.
6. Z. Liu, H. Sun, G. Marine, and H. Wu, "6G IoV networks driven by RF digital twin modeling," IEEE Trans. Intell. Transp. Syst., vol. 25, no. 3, pp. 2976-2986, Mar. 2024, doi: 10.1109/TITS.2023.3292533.
7. P. Yin, W. Liang, J. Wen, J. Kang, J. Chen, and D. Niyato, "Multi-agent DRL for multi-objective twin migration routing with workload prediction in 6G-enabled IoV," arXiv preprint arXiv:2505.07290v2, Mar. 2026.
8. J. Zeng, Z. Shi, C. Li, M. Yan, H. Zhang, S. Chen, X. Hu, and X. Li, "Digital twin-enabled mobility-aware cooperative caching in vehicular edge computing," arXiv preprint arXiv:2603.06653v1, Feb. 2026.
9. S. Zhan, Z. Liu, L. Huang, S. Shen, Z. Bai, Z. Gao, and D. Niyato, "Digital twin-assisted task offloading and resource allocation in ISAC-enabled Internet of Vehicles," arXiv preprint arXiv:2511.05789v1, Nov. 2025.
10. H. Hu, Z. Song, and W. Shi, "Large AI model-enhanced digital twin-driven 6G healthcare IoE," Electronics, vol. 15, art. no. 00619, 2026.
11. A. Zaki-Hindi et al., "A reference functional architecture for network digital twins in 6G systems," IEEE Open J. Commun. Soc., 2026, doi: 10.1109/OJCOMS.2026.3668035.
12. L. Li, "Research on digital twin technology for future 6G," Intelligent and Converged Networks, vol. 6, no. 1, pp. 20-40, 2025, doi: 10.23919/ICN.2025.0002.
13. S. J. Siddiqi, A. H. Alobaidi, M. A. Jan, and M. Tariq, "Securing vehicle-to-digital twin communications in the Internet of Vehicles," ACM Trans. Internet Technol., 2025, doi: 10.1145/3711863.
14. V. Farre, J. Vega-Sanchez, V. H. Garzon Pacheco, N. V. Orozco Garzon, J. C. Estrada-Jimenez, J. A. Vasquez-Peralvo, and S. Chatzinotas, "Digital twin-driven cross-layer orchestration using multi-agent reinforcement learning for 6G UAV-NTN networks," IEEE Access, 2026, doi: 10.1109/ACCESS.2026.3690691.
15. J. Wang, K. Zhu, and E. Hossain, "Green Internet of Vehicles (IoV) in the 6G era: Toward sustainable vehicular communications and networking," IEEE Trans. Green Commun. Netw., vol. 6, no. 1, pp. 391-423, Mar. 2022.
16. X. Yuan, J. Chen, N. Zhang, J. Ni, F. R. Yu, and V. C. M. Leung, "Digital twin-driven vehicular task offloading and IRS configuration in the Internet of Vehicles," IEEE Trans. Intell. Transp. Syst., vol. 23, no. 12, pp. 24290-24304, Dec. 2022.
17. Y. Hui, X. Ma, C. Li, N. Cheng, R. Chen, Z. Yin, T. H. Luan, and G. Mao, "Service-oriented edge collaboration: Digital twin enabled edge collaboration for composite services in AVNs," IEEE Trans. Intell. Transp. Syst., vol. 26, no. 3, pp. 3147-3160, Mar. 2025.
18. L. Tang, X. Liao, J. Yang, L. Wang, Q. Liu, and Q. Chen, "Digital twin-based joint optimization strategy for dual time-domain slice resource management and DT deployment in IoV," IEEE Internet of Things J., vol. 12, no. 24, pp. 52748-52762, Dec. 2025.
19. X. Wang, C. He, W. Jiang, W. Wang, L. Li, and X. Xie, "UAV-assisted task offloading and resource allocation in Internet of Vehicles: An integration of digital twin and generative AI," IEEE Trans. Netw. Sci. Eng., vol. 13, pp. 5038-5055, 2026.
20. J. A. Ansere, S. C. Prabhashana, O. A. Dobre, and T. Q. Duong, "Quantum machine learning DDPG for digital twin semantic vehicular networks," in Proc. IEEE Int. Conf. Machine Learning for Communications and Networking (ICMLCN), 2025.
21. A. K. Selvaraj, Y. Govindarajan, S. Beni Prathiba, A. Aashish Vinod, V. P. Amirtha Ganesan, Z. Zhu, and T. R. Gadekallu, "Federated learning and digital twin-enabled distributed intelligence framework for 6G autonomous transport systems," IEEE Trans. Intell. Transp. Syst., vol. 26, no. 10, pp. 18214-18227, Oct. 2025.

DOI: http://doi.org/10.5281/zenodo.20557590

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

22. 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 Internet of Things J., vol. 13, no. 10, pp. 19861-19876, May 2026.
23. L. U. Khan, W. Saad, D. Niyato, Z. Han, and C. S. Hong, "Digital-twin-enabled 6G: Vision, architectural trends, and future directions," IEEE Commun. Mag., vol. 60, no. 1, pp. 74-80, Jan. 2022.
24. W. Saad, M. Bennis, and M. Chen, "A vision of 6G wireless systems: Applications, trends, technologies, and open research problems," IEEE Network, vol. 34, no. 3, pp. 134-142, May 2020.
25. K. B. Letaief, W. Chen, Y. Shi, J. Zhang, and Y. A. Zhang, "The roadmap to 6G: AI empowered wireless networks," IEEE Commun. Mag., vol. 57, no. 8, pp. 84-90, Aug. 2019.
26. H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, 2017.
27. E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in Proc. ICLR, 2022.

DOI: http://doi.org/10.5281/zenodo.20557590