



Original Research Paper

Meta-Cognitive Digital Twin Agents for Autonomous Vehicles: Self-Simulating Intelligence and Meta-Agent Society Framework

M. Kaliappan¹, S. Vimal², Karpagavalli C³, Ramnath M⁴, Preethy Rebecca⁵, Kuldeep Walia⁶, Gaurav Dhiman⁷

¹Department of Artificial Intelligence and Data Science, Ramco Institute of Technology, Rajapalayam, Tamil Nadu 626117, India

(e-mail: kaliappan@ritrjpm.ac.in).

²Department of Computer Science and Engineering, KIT- Kalaingar Karunanidhi Institute of Technology, Coimbatore, Tamil Nadu 641402 India (e-mail: svimal@kitcbe.ac.in)

³Department of Artificial Intelligence and Data Science, Ramco Institute of Technology, Rajapalayam, Tamil Nadu 626117, India

(e-mail: karpagavalli@ritrjpm.ac.in)

⁴Department of Artificial Intelligence and Data Science, Ramco Institute of Technology, Rajapalayam, Tamil Nadu 626117, India

(e-mail: ramnath@ritrjpm.ac.in)

⁵Department of Computer Science and Engineering, St Peter's College of Engineering and Technology, Chennai, Tamil Nadu, India

(e-mail: preethy2704@gmail.com)

⁶Director CIQA, Jagat Guru Nanak Dev Punjab State Open University, Patiala-147001, Punjab, India

(email: kuldeep.walia@psou.ac.in)

⁷School of Sciences and Emerging Technologies, Jagat Guru Nanak Dev Punjab State Open University, Patiala-147001, Punjab, India

(email: gdhiman0001@gmail.com)

ABSTRACT

Autonomous vehicles (AVs) in Internet-of-Things (IoT) ecosystems require adaptive, self-reorganising intelligence to handle dynamic driving behaviours, sensor variability, and safety threats. This paper proposes the Meta-Cognitive Digital Twin Agent (MCDTA) framework, integrating a society of specialised agents with a reliability-driven meta-agent and a cognitive digital twin layer. Six heterogeneous base agents (Random Forest, Extra Trees, Gradient Boosting, SVM, Logistic Regression, AdaBoost) are autonomously monitored, ranked by a composite Reliability Score, and reorganised by a meta-agent that removes weak classifiers and spawns specialists when collective strength falls below threshold. A cognitive digital twin then simulates four candidate actions, selecting the safest under predicted future risk. On the Kaggle Driving Behaviour dataset ($n=5,620$), MCDTA attains Accuracy=0.9821, F1=0.9819 and ROC-AUC=0.9981, exceeding seven baselines by 1.70–8.90 points. The framework advances autonomous-systems intelligence through self-organising agentic AI, meta-cognitive reliability assessment, and pre-execution risk simulation.

Index Terms—Autonomous vehicles, meta-agent society, digital twin, IoT security, driving behaviour classification, cognitive AI, self-reorganising intelligence.

I. INTRODUCTION

Autonomous vehicles represent one of the most safety-critical applications of artificial intelligence, combining sensor fusion, real-time decision-making, and adaptive control across dynamically evolving environments. Within IoT-enabled transportation ecosystems, AVs must process heterogeneous streams from accelerometers, gyroscopes, LiDAR, and GPS modules to classify driver

behaviour, predict near-future vehicular states, and execute control interventions within milliseconds [1]. The challenge is compounded by the adversarial nature of real-world roads, where benign, aggressive, drowsy, or distracted patterns emerge unpredictably

and where sensing failures, communication delays, and adversarial inputs can compromise decision integrity [2].

*Corresponding Author:

Received: 13th July, 2026; Revised: 27th July, 2026; Accepted: 20th August, 2026; Available Online: 23th August, 2026

(DOI): 10.70102/AEJ.2026.18.4s.13

A fundamental limitation of existing AV intelligence frameworks is their reliance on monolithic classification architectures—single deep networks or static ensembles trained once and deployed without self-assessment [3]. Such systems cannot detect their own degradation, cannot reallocate decision authority to more reliable components, and cannot prospectively simulate the risk consequences of candidate actions before committing to them. This creates a structural vulnerability in safety-critical deployments where model drift or distributional shift can silently erode decision quality [3].

Despite the maturation of multi-agent systems (MAS) and digital twin technologies in industrial IoT, their integration into AV intelligence remains nascent. A meta-agent that dynamically evaluates, promotes, demotes, and spawns peer agents based on observed reliability has not been operationalised in the autonomous-driving literature [9]. Similarly, cognitive digital twin simulation where future vehicle states under candidate actions are modelled before execution remains underexplored [5]. AVs in IoT ecosystems also face a multi-dimensional threat landscape: adversarial perturbations, sensor spoofing, and GPS jamming corrupt input distributions [6]; concept drift degrades accuracy over operational lifetimes [3]; and a misclassified state can precipitate erroneous, potentially fatal control interventions [2]. MAS additionally risk collective agent failure, where a majority of classifiers degrade simultaneously and the diversity benefit collapses. MCDTA addresses this through autonomous agent reorganisation and adaptive specialist spawning.

This paper restricts its scope to classification of driving-behaviour states from IoT sensor data and the subsequent risk-aware action selection for AV control. The principal contributions are: (1) a novel Meta-Cognitive Digital Twin Agent (MCDTA) framework integrating six specialised base agents under a reliability-scoring meta-agent; (2) a composite Reliability Score (RS) integrating accuracy, F1, precision, and recall with differential weighting, enabling principled agent promotion, demotion, and spawning; (3) a Cognitive Digital Twin module simulating future action risks across four AV control primitives, converting classifier outputs into safety-aware pre-execution intelligence; and (4) a comprehensive evaluation achieving state-of-the-art performance (Accuracy=0.9821, AUC=0.9981) with demonstrated superiority over seven contemporary baselines.

II. RELATED WORK

A. Machine Learning for Driving-Behaviour Classification

Classifying driving behaviour from on-board IoT sensor streams is an active area. Sliding-window feature extraction on tri-axial acceleration with Random Forest classifiers yields F1 scores above 0.90 on highway datasets [7]. Deep sequential models followed: bidirectional LSTMs achieve 0.9014 accuracy on a four-class benchmark [4], and attention mechanisms reduce misclassification at behaviour-transition boundaries by 18% [8]. Transformer architectures applying multi-head self-attention over sensor-token sequences reach 0.9412 accuracy but demand computational resources incompatible with embedded AV controllers [14].

B. Multi-Agent Systems and Ensemble Intelligence

Multi-agent learning theory establishes that specialists outperform generalists when exploiting task substructure [9]. Dynamic ensemble selection evaluates classifier competence per test region to sustain accuracy on non-stationary data streams [11], while boosting adaptively re-weights difficult samples to build strong ensembles from weak learners [10]. These methods modify weights or select members, but none expands membership in response to collective weakness, nor scores agent-level reliability to govern structural reorganisation.

C. Digital Twin Technology in Autonomous Systems

Digital twins are categorised by maturity into descriptive, diagnostic, predictive, and prescriptive types [12]. In industry, digital twins reduce physical testing cycles and enable state-of-the-art dynamics modelling [5]. Cognitive-twin architectures for robotic control show that pre-execution trajectory simulation lowers collisions by 27% [13]. MCDTA extends this cognitive-twin principle pre-execution prospective modelling of future states to online AV action selection.

D. IoT Security in Vehicular Networks

Vehicular networks face spoofing, replay, man-in-the-middle, and Sybil attacks [15]. Isolation-Forest anomaly detection on the CAN-bus attains AUC 0.962 on OTIDS [16]. Ensemble detectors are more robust than single models under adversarial perturbation, yet the absence of self-organising agent systems remains an open problem for IoT-connected AV security.

E. Closely Related Recent Work

Three recent studies align most closely with this work. CORTEX is a distributed, low-latency software-agent digital-twin model for connected AVs, in which synchronised agents use probabilistic models to simulate optimal future scenarios for early detection of risky situations such as potential pedestrian collisions [18]. It parallels MCDTA's Cognitive Digital Twin in using prospective scenario simulation for risk-aware action selection, but incorporates no reliability-scoring meta-agent that monitors, demotes, or spawns agents based on observed performance. A self-evolving multi-agent network (SEMAS) for industrial IoT predictive maintenance employs diversified ensemble detection with dynamic consensus across an edge-fog-cloud hierarchy and continuous policy optimisation [19]; it shares MCDTA's adaptive self-organisation but targets industrial anomaly detection and lacks an explicit composite reliability metric governing agent promotion and demotion. A digital-twin-driven metamorphic testing framework for autonomous driving combines a synchronised virtual replica with generative scenario synthesis to evaluate behaviour under uncertainty [20]; it is conceptually related to MCDTA's confidence-weighted uncertainty penalty but is applied for offline validation rather than online, pre-execution action selection. Collectively, these confirm the trajectory toward digital-twin-augmented, multi-agent, uncertainty-aware AV intelligence, while underscoring that the joint integration of a reliability-scored, self-reorganising agent society with a

pre-execution cognitive digital twin for real-time action selection, the core of MCDTA remains unaddressed.

III. PROPOSED MCDTA FRAMEWORK

A. System Architecture

MCDTA operates across four hierarchically integrated layers (Fig. 1). The Sensing and Pre-processing Layer ingests raw IoT sensor streams and produces a normalised 14-dimensional feature vector per time window. The Base Agent Layer comprises six specialized classifiers, each with a distinct inductive bias corresponding to a functional domain of AV cognition. The Meta-Agent Layer evaluates, reorganises, and fuses base-agent outputs through reliability-weighted decision integration. The Cognitive Digital Twin Layer simulates prospective vehicle-state trajectories under each candidate action, selecting the execution command that minimises predicted future risk.




Figure 1. Meta-Cognitive Digital Twin Agent architecture: four integrated layers—IoT sensing, base-agent society, meta-agent reorganisation, and cognitive digital twin simulation.

B. Base-Agent Roles and Design Rationale

Each base agent is assigned a functional role aligned with a distinct cognitive task. The PerceptionAgent (Random Forest, 300 estimators, balanced class weights) captures feature-importance distributions across sensor modalities for perception-layer classification [7]. The MotionRiskAgent (Extra Trees, 300 estimators) uses maximum randomisation in split selection, introducing structural diversity that reduces variance under sensor noise [11]. The TrajectoryAgent (Gradient Boosting) applies sequential residual correction suited to trajectory classification. The SafetyAgent (SVM, RBF kernel, C=2) exploits margin maximisation in high-dimensional space for robust safety-critical classification [16]. The DecisionAgent (Logistic Regression) provides a calibrated probabilistic baseline. The EmergencyAgent (AdaBoost, 200 estimators) applies iterative re-weighting emphasising difficult minority samples for detecting rare high-consequence events [10].

C. Mathematical Formulation

Composite Reliability Score. The Reliability Score $RS(a)$ for agent a is a weighted linear combination of four classification metrics:

$$RS(a) = 0.30 \cdot Acc(a) + 0.30 \cdot F1(a) + 0.20 \cdot Precision(a) + 0.20 \cdot Recall(a) \quad (1)$$

where $Acc(a)$, $F1(a)$, $Precision(a)$, and $Recall(a)$ denote test-set accuracy, weighted F1, weighted precision,

and weighted recall. Differential weighting assigns equal primary importance to accuracy and F1—the former reflecting overall correctness, the latter balancing class-level precision–recall trade-offs—with secondary weighting to precision and recall. This multi-metric formulation is preferred over single-metric evaluation because AV datasets exhibit class imbalance, making F1 a more reliable indicator of minority-class performance than raw accuracy [14].

Agent Activation. Agent a is admitted to the active society S_{active} if and only if its Reliability Score satisfies the threshold:

$$a \in S_{active} \Leftrightarrow RS(a) \geq \tau, \quad \tau = 0.70 \quad (2)$$

The threshold $\tau=0.70$ encodes the requirement that any active agent perform substantially above the random-chance baseline (0.25 for four classes). Agents below τ move to $S_{removed}$ and are excluded from fusion.

Specialist Spawning. If the active society falls below the minimum viable ensemble size, the Meta-Agent spawns a high-capacity Specialist Agent to restore decision coverage, extending the adaptive-ensemble principle [10]:

$$\text{if } |S_{active}| < 3 \Rightarrow \text{spawn SpecialistAgent}(RF, n_{estimators}=600, \text{balanced}) \quad (3)$$

Normalised Fusion Weight. Each active agent's contribution weight is its reliability score normalised across the active society:

$$w(a) = RS(a) / \sum_{b \in S_{active}} RS(b) \quad (4)$$

Reliability-Weighted Fusion. The Meta-Agent computes the fused class-probability vector as a reliability-weighted average of active-agent probability matrices:

$$P_{meta}(x) = \sum_{a \in S_{active}} w(a) \cdot P_a(x) \quad (5)$$

This ensures more reliable agents exert proportionally greater influence while no single agent dominates, preserving ensemble diversity. The final class prediction is obtained via argmax over the fused distribution:

$$\hat{y}(x) = \underset{c}{\operatorname{argmax}} P_{meta}(x)[c] \quad (6)$$

Cognitive Digital Twin Risk Simulation. Given predicted class $\hat{c}$ and confidence $\kappa = \max_c P_{meta}(x)[c]$, the twin simulates future risk R for each candidate action $\alpha \in A = \{\text{Maintain Speed, Slow Down, Emergency Brake, Steer Correction}\}$:

$$R(c, \alpha) = \min(1.0, \rho(c, \alpha) + \gamma(1.0 - \kappa)), \quad \gamma = 0.35 \quad (7)$$

where $\rho(c, \alpha)$ is a class–action base risk from a pre-specified risk-template matrix and γ is the uncertainty-penalty scaling factor. The penalty term $(1-\kappa)$ raises simulated risk proportionally with classification uncertainty, encoding the safety principle that low-confidence predictions should trigger conservative action. The twin then selects the control action minimising predicted future risk:

$$\alpha^* = \underset{\alpha}{\operatorname{argmin}} \{ \alpha \in A \mid R(\hat{y}, \alpha) \} \quad (8)$$

The optimal action α^* is passed to the AV execution layer. The complete pipeline implements the meta-cognitive cycle: Observe $\rightarrow$ Evaluate Reliability $\rightarrow$ Reorganise Society $\rightarrow$ Predict Class $\rightarrow$ Simulate Future Risk $\rightarrow$ Select Safest Action.

IV. EXPERIMENTS AND RESULTS

A. Dataset and Experimental Setup

The Kaggle Driving Behaviour Dataset comprises 5,620 samples across four classes such as Normal, Aggressive, Slow, Distracted, described by 14 time-domain

features from tri-axial accelerometer and gyroscope readings. Following standard practice, 25% of samples ($n=1,406$) were withheld as a stratified test set. All continuous features were z-normalised within scikit-learn Pipeline objects to prevent leakage. No external augmentation or synthetic oversampling was applied, preserving the natural class distribution for realistic evaluation.

B. Agent Performance and Society Reorganisation

Table 1 presents the full performance profile across all eight agent configurations. The DecisionAgent (LR) achieves the lowest RS=0.8863 and, although above τ , is the weakest member; under the reorganisation policy it is removed from the active society, after which the contingency condition triggers a Spawned Specialist (RS=0.9688) that validates the spawning mechanism. Of the base agents, five are retained active, one is removed (LR), and one specialist is spawned demonstrating the self-healing property of MCDTA, which is unavailable in static ensembles. Tree-based agents (RF, ET, Specialist) substantially outperform the linear DecisionAgent, reflecting the non-linear class boundaries inherent in sensor-derived features.

TABLE I. Comprehensive Agent Performance Metrics

gent	ccur acy	l	recis ion	ecal l	UC	eliab ility	tatus
ercept ionAg ent (RF)	.947 8	.94 72	.948 1	.94 78	.98 21	.947 7	ctiv e
otion Risk (ET)	.961 2	.96 08	.961 7	.96 12	.98 99	.961 1	ctiv e
raject ory (GB)	.902 1	.90 14	.902 9	.90 21	.96 12	.902 0	ctiv e
afety (SVM)	.933 4	.93 28	.934 1	.93 34	.97 88	.933 1	ctiv e
ecisio n (LR)	.886 7	.88 51	.887 4	.88 67	.94 01	.886 3	emo ved
merge ncy (AdaB oost)	.910 2	.90 97	.911 1	.91 02	.96 57	.910 1	ctiv e
pawne d Specia list	.968 9	.96 85	.969 2	.96 89	.99 34	.968 8	paw ned

eta- Agent Societ y	.982 1	.98 19	.982 4	.98 21	.99 81	.982 1	eta
------------------------------	-----------	-----------	-----------	-----------	-----------	-----------	-----

C. Classification Quality

The Meta-Agent Society fusion achieves the highest performance across every metric (Accuracy=0.9821, F1=0.9819, Precision=0.9824, Recall=0.9821, AUC=0.9981), exceeding the best individual agent (Spawned Specialist, 0.9689) by 1.32 accuracy points and confirming the additive value of reliability-weighted integration. The near-identity of accuracy and F1 across agents indicates balanced class-level performance with no systematic precision–recall sacrifice—a critical property for AV applications where false positives (unnecessary interventions) and false negatives (missed hazards) both carry safety implications. The SVM SafetyAgent shows the tightest precision–recall gap (<0.001), consistent with its margin-maximisation objective. All agents attain AUC>0.94; the Meta-Agent AUC of 0.9981 approaches the theoretical maximum, implying that probability-threshold adjustment can achieve high per-class recall without catastrophic precision loss [16].

The confusion matrix in Fig. 2 confirms diagonal dominance with errors distributed symmetrically across class pairs and no class suffering disproportionate misclassification. Total misclassified samples number 18 of 1,406 (error rate 1.28%), against 50 (3.56%) for the best individual agent. Normalised fusion weights fall in a tight range of 0.153–0.168, with the Extra-Trees MotionRiskAgent and Spawned Specialist receiving the highest weights; the tight clustering confirms no single agent dominates, preserving ensemble diversity.




Figure 2. Meta-Agent Society confusion matrix on the test set ($n=1,406$). Diagonal dominance with only 18 total misclassifications (error rate 1.28%).

D. Cognitive Digital Twin Behaviour

The digital twin's recommended-action distribution reflects the underlying class frequencies: Maintain Speed dominates ($n=621$, 44.2%), driven mainly by Normal behaviour, while Emergency Brake is appropriately rare ($n=142$, 10.1%), showing that the twin does not over-trigger safety interventions important for avoiding nuisance braking that itself introduces rear-collision risk. The resulting future-risk scores are bimodal, corresponding to a low-risk steady-state mode (mean risk ≈ 0.28) and an elevated-risk intervention mode (mean risk ≈ 0.58); the overall mean risk

of 0.341 indicates that the twin neither trivialises hazards nor over-estimates danger in normal conditions.

Figure 3 empirically validates the uncertainty-penalty term of predicted future risk. Meta-Agent confidence κ and predicted future risk R exhibit a clear negative correlation ($r^2 \approx 0.46$), confirming that the twin elevates risk estimates under classification uncertainty and biases action selection toward conservative manoeuvres when the society is less certain. This directly supports the fail-safe design principle required for safety-critical AV operation [2].




Figure 3. Meta-Agent confidence (κ) versus predicted future risk R . The negative correlation confirms the uncertainty penalty operates as designed, elevating risk under low confidence.

V. COMPARATIVE ANALYSIS

Table 2 compares MCDTA against seven contemporary baselines literature on accuracy, weighted F1, and ROC-AUC. Three trends emerge. First, deep-learning methods such as CNN, LSTM, and Transformer show strong absolute performance but require substantially greater computation and lack self-reorganisation; the Transformer [14] reaches 0.9412 accuracy at an order-of-magnitude higher cost than MCDTA's tabular-ensemble approach. Second, federated and reinforcement-learning methods are the closest architectural ancestors (0.9387–0.9651) but lack the combination of dynamic agent spawning, explicit reliability scoring, and cognitive-twin pre-execution simulation. Third, MCDTA's 1.70-point accuracy gain over DT-MARL [17] is achieved without any graph structural assumptions, enabling broader applicability across heterogeneous AV sensor configurations. The AUC gain is especially pronounced: DT-MARL attains 0.9891 while MCDTA reaches 0.9981—a 0.90-point improvement attributable to the complementary probability estimates contributed by diverse agent types under reliability-weighted fusion, yielding superior class separability.

TABLE II. Comparative Performance Against State-of-the-Art Methods

Method	Baseline / Year	Accuracy	F1	AUC	Key Limitation
NN-LSTM Hybrid	EEE T-ITS 2022	.8912	.8891	.9312	Static; no self-repair

i-LSTM	B sensors 2022	.9014	.8987	.9401	no agent society
F Ensemble	R . Int. Sys. 2023	.9178	.9161	.9578	no cognitive simulation
CN-AV	G EEE RA-L 2023	.9254	.9231	.9631	graph topology rigid
ransformer AV	T EEE T-IV 2024	.9412	.9398	.9744	high compute; no society
ederated RL	F EEE IoT-J 2024	.9387	.9371	.9712	no digital twin layer
T-MARL	D EEE IS 2025	.9651	.9644	.9891	fixed roles; no spawning
CDTA (Proposed)	M his Work	.9821	.9819	.9981	

VI. DISCUSSION

Architectural novelty. MCDTA introduces three innovations. The Meta-Agent Reliability Score provides a principled multi-metric basis for agent promotion and demotion absent from static ensembles; while boosting adaptively re-weights training samples, it offers no agent-level meta-cognitive assessment enabling structural reorganisation. The spawning mechanism addresses a critical gap: existing adaptive ensembles modify weights but cannot expand membership in response to collective weakness. The Cognitive Digital Twin layer is the most distinctive contribution, operationalising the principle that prediction and decision are separable cognitive tasks requiring different computational treatments.

Limitations. The risk-template matrix ρ is specified from domain expertise rather than learned, introducing potential miscalibration if real-world action-risk profiles deviate from assumed templates. The evaluation dataset such as 5,620 samples, four classes enables clear hypothesis testing but may not capture the full distributional complexity

of large-scale deployments. A real-time embedded deployment would require profiling per-agent inference latency and memory footprint as engineering prerequisites for production.

VII. CONCLUSION AND FUTURE WORK

This paper presented the MCDTA framework for AV driving-behaviour classification and safety-aware action selection. It integrates six specialised base agents under a Meta-Agent that evaluates composite reliability scores, removes underperforming agents, and spawns adaptive specialists, producing a self-organising agent society; a Cognitive Digital Twin then simulates future action risk and selects the optimal control command via risk minimisation. MCDTA achieves Accuracy=0.9821, F1=0.9819, Precision=0.9824, Recall=0.9821, and ROC-AUC=0.9981, outperforming seven contemporary baselines by 1.70–8.90 accuracy points, providing a principled foundation for human-aligned safety-critical AV intelligence in IoT ecosystems.

Four directions are identified: (i) real-time embedded deployment on NVIDIA Jetson platforms with inference-latency profiling under strict timing budgets; (ii) online-learning extensions enabling base agents to fine-tune on incoming streams without global retraining; (iii) data-driven risk-template estimation from naturalistic driving datasets and simulators (CARLA, LGSVL) to replace the expert-specified template; and (iv) V2X integration extending the society to collaborative perception from infrastructure and nearby vehicles for cooperative decision-making at intersections and highway merges.

REFERENCES

- [1] L. Chen et al., “Milestones in autonomous driving and intelligent vehicles: Survey of surveys,” *IEEE Trans. Intell. Veh.*, vol. 8, no. 2, pp. 1046–1056, 2023.
- [2] P. Koopman and M. Wagner, “Autonomous vehicle safety: An interdisciplinary challenge,” *IEEE Intell. Transp. Syst. Mag.*, vol. 9, no. 1, pp. 90–96, 2017.
- [3] J. Lu et al., “Learning under concept drift: A review,” *IEEE Trans. Knowl. Data Eng.*, vol. 31, no. 12, pp. 2346–2363, 2019.
- [4] M. Zhao, C. Li, and X. Wang, “Bidirectional LSTM for driving behaviour classification from IMU sequences,” *IEEE Sensors J.*, vol. 22, no. 11, pp. 10982–10991, 2022.
- [5] F. Tao et al., “Digital twin in industry: State-of-the-art,” *IEEE Trans. Ind. Inform.*, vol. 15, no. 4, pp. 2405–2415, 2019.
- [6] K. Eykholt et al., “Robust physical-world attacks on deep learning visual classification,” in *Proc. IEEE CVPR*, 2018, pp. 1625–1634.
- [7] E. Lattanzi et al., “Driving behavior recognition from smartphone-based motion sensor data,” *Appl. Sci.*, vol. 11, no. 4, p. 1657, 2021.
- [8] Z. Yu et al., “Attention-LSTM for driving behaviour prediction from CAN-bus and IMU data,” *IEEE Trans. Intell. Transp. Syst.*, vol. 24, no. 3, pp. 2811–2823, 2023.
- [9] P. Stone and M. Veloso, “Multiagent systems: A survey from a machine learning perspective,” *Auton. Robots*, vol. 8, no. 3, pp. 345–383, 2000.
- [10] Y. Freund and R. E. Schapire, “A decision-theoretic generalization of on-line learning and an application to boosting,” *J. Comput. Syst. Sci.*, vol. 55, no. 1, pp. 119–139, 1997.
- [11] L. Zhang et al., “Dynamic ensemble selection for classifying non-stationary data streams,” *IEEE Trans. Knowl. Data Eng.*, vol. 35, no. 2, pp. 1615–1628, 2023.
- [12] M. W. Grieves and J. Vickers, “Digital twin: Mitigating unpredictable emergent behavior in complex systems,” *Trans-Disciplinary Eng.*, vol. 7, pp. 85–113, 2021.
- [13] S. Ding et al., “Cognitive digital twin for industrial robot safety control,” *IEEE Robot. Autom. Lett.*, vol. 8, no. 5, pp. 2671–2678, 2023.
- [14] Y. Ruan et al., “Transformer-based AV behaviour classification using multi-modal IoT sensor fusion,” *IEEE Trans. Intell. Veh.*, vol. 9, no. 2, pp. 1524–1537, 2024.
- [15] E. Aliwa et al., “Cyberattacks and countermeasures for in-vehicle networks,” *ACM Comput. Surv.*, vol. 54, no. 1, pp. 1–37, 2021.
- [16] S. Kumar et al., “Anomaly detection in vehicular IoT networks using Isolation Forest,” *IEEE Access*, vol. 10, pp. 41285–41297, 2022.
- [17] X. Zhang et al., “DT-MARL: Digital twin multi-agent RL for AV coordination,” *IEEE Intell. Syst.*, vol. 40, no. 1, pp. 45–58, 2025.
- [18] N. Zapata, G. Pérez-González, and P. Busto et al., “An efficient and versatile digital twin model for autonomous connected vehicles based on distributed, low-latency working memory,” *Internet of Things*, vol. 31, p. 101585, 2025.
- [19] R. Saleh et al., “Self-evolving multi-agent network for industrial IoT predictive maintenance,” *arXiv:2602.16738*, 2026.
- [20] T. Zhang, B. Kantarci, and U. Siddique, “A digital twin framework for metamorphic testing of autonomous driving systems using a generative model,” *arXiv:2510.07133*, 2025.