UNSYS digital

The Journal of Instrumentation, Automation, and System

Multi-Agent Simulation for Coordinated Aerial Surveillance using Digital Twin and SITL Framework

Agus Budiyono[†], Ivan[✓], Vishnu K. Kaliappan[♪]

[†]CAP Solutions, Jakarta, Indonesia.
[✓]Holoxit.
[♪]KPR IET, Coimbatore, India.

Abstract— The increasing complexity of aerial surveillance tasks in urban and industrial environments demands scalable, coordinated, and resilient multi-UAV systems. This paper presents a Digital Twin–Software-in-the-Loop (DT–SITL) framework for learning-based cooperative aerial surveillance, enabling systematic development and validation of multi-drone tracking behaviors prior to physical deployment. Within the proposed framework, each UAV is represented by a synchronized Digital Twin instance and controlled by a shared Proximal Policy Optimization (PPO) policy, allowing decentralized execution while preserving coordinated system-level behavior. The framework integrates UAV dynamics, target motion modeling, structured observation construction, and reward-driven multi-agent learning into a unified DT–SITL execution loop. Simulation results involving three UAVs tracking a dynamically maneuvering target demonstrate stable policy convergence, bounded tracking error, and emergent cooperative behaviors such as spatial distribution and persistent engagement. System-level analysis shows that the DT–SITL approach enables safe experimentation, repeatable evaluation, and scalable coordination analysis, establishing a practical foundation for future autonomous multi-UAV surveillance systems and their transition toward real-world deployment.

Keywords— digital twin, software-in-the-loop simulation, UAV

Copyright© 2025. Published by UNSYSdigital. All rights reserved.
DOI: 10.21535/36bqby95

I. Introduction

URBAN security operations increasingly demand real-time, distributed sensing to monitor complex and dynamic environments. Traditional fixed surveillance systems often lack adaptability, and single-UAV surveillance missions struggle to maintain persistent situational awareness in large or occlusion-prone areas. Multi-UAV systems offer scalable coverage and flexible mission execution but introduce challenges in decentralized coordination, reliable communication, cooperative perception, and safety assurance during operation.

Digital Twin (DT) technology provides a compelling opportunity to address these challenges. A DT enables a continuously updated virtual representation of physical UAV assets, supporting predictive analytics, mission validation, and autonomous decision-making. In the UAV field, DTs have been used to model radio communication (Aláez et al., 2024), support AI-assisted operations (Gürses et al., 2024), and enable secure federated learning pipelines (Zami, Uddin, and Nguyen, 2025). When coupled with Software-in-the-Loop (SITL) testing, DTs allow realistic verification of autonomy and communication logic without requiring live flight operations.

This paper presents a Digital Twin–Software-in-the-Loop (DT-SITL) framework for multi-UAV coordination. Built on AirSim, the framework supports multi-agent simulation where UAVs collaboratively perform surveillance tasks such as object tracking, area coverage, and dynamic target following. The system provides three main contributions:

1. A scalable DT-SITL architecture for multi-agent UAV simulation, enabling synchronized DT states, realistic sensor modeling, and controller integration.
2. A cooperative tasking and tracking framework supported by inter-UAV communication and distributed decision-making.
3. Extensive simulation validation demonstrating cooperative tracking behavior, optimized coverage, and safe navigation in dynamic urban environments.

II. Literature Review

A. *Multi-Agent UAV Coordination*

Multi-agent UAV coordination has emerged as a pivotal research area, driven by increasing operational requirements for large-scale, persistent, and adaptive aerial surveillance. Coordination enables UAVs to function not as isolated platforms but as distributed autonomous agents capable of sharing information, negotiating task assignments, and maintaining global situational awareness. This paradigm is supported by graph-based models, distributed optimization, consensus algorithms, and multi-agent reinforcement learning (MARL), depending on mission complexity and communication constraints.

Early foundational studies conceptualized multi-UAV surveillance through formal spatio-temporal specifications, in which UAV trajectories were generated to satisfy coverage and

Corresponding author: Agus Budiyono (e-mail: budiyono@alum.mit.edu)

timing requirements (Ahmadzadeh et al., 2006). These works demonstrated the feasibility of encoding mission constraints mathematically but relied on centralized planners, limiting scalability. Subsequent efforts shifted to decentralized coordination, enabling UAVs to collaboratively plan and adapt without reliance on a single command node (Gray et al., 2025).

More recently, MARL approaches have achieved substantial advances by allowing UAV teams to learn cooperative behaviors through shared rewards or multi-objective policies. Su and Qian (2023) and Ekechi et al. (2025) discuss the ability of MARL agents to learn robust strategies for tracking, search-and-rescue, and coverage in dynamic environments. Under MARL, UAVs model both the environment and the actions of their teammates, capturing complex interactions such as formation maintenance, cooperative pursuit, and obstacle-aware negotiation.

Communication plays a central role in multi-agent coordination. Fault-tolerant communication topologies—such as mesh networks, relay-assisted connectivity, or priority-aware links—help ensure task execution even under partial failures (Yahya et al., 2025). Emerging concepts such as quality-of-service (QoS) guarantees (Seraj et al., 2022) also enhance reliability for time-critical missions like wildfire monitoring or multi-target tracking.

Overall, the literature shows a transition from centralized trajectory assignment → distributed consensus → learning-based cooperative intelligence, forming the methodological basis upon which Digital Twin (DT) and SITL frameworks integrate multi-agent UAV behaviors in virtual environments.

B. Digital Twin Framework in UAV Systems

Digital Twin (DT) technology is transforming UAV research by enabling real-time bidirectional coupling between virtual models and physical assets. For UAVs, the DT paradigm provides high-fidelity mirroring of flight dynamics, environmental interactions, sensor states, hardware configurations, and mission execution logic. This allows developers to test mission plans, failure cases, and cooperative behaviors in the digital environment before physical execution, reducing cost and risk.

DTs for UAV systems generally involve three components:

1. The physical UAV
2. A virtual UAV model with equivalent dynamics, controllers, and sensors
3. A synchronization mechanism for state, environment data, and mission updates

Danish Iqbal & Buhnova (2024) outline the architectural requirements for UAV DT designs, emphasizing modular modeling, continuous synchronization, and reliability. McClellan et al. (2021) demonstrate a physics-based DT for UAV landing control, where online model predictive control (MPC) leverages digital state estimation for precision landing, a powerful demonstration of DT-enabled autonomy.

A growing trend is the integration of DTs with large-scale autonomous networks. Studies such as Gürses et al. (2024) and Zhou, Leng & Wang (2024) show how DTs support AI research for UAV swarms, enabling:

- scenario replay and forensic analysis
- hardware-in-the-loop (HIL) and SITL co-simulation
- large-scale virtual testing for network protocols and cooperative perception
- federated learning for distributed UAV intelligence

DTs enhance system reliability by providing a risk-free yet physically grounded test environment. Radio communication modeling using DTs, as explored by Aláez et al. (2024), further supports UAV operations in complex terrains and urban areas.

In swarm missions, multi-UAV DTs enable continuous monitoring of fleet health, distributed diagnostics, and coordinated responses to anomalies. They also provide a platform for assessing emergent behaviors (such as flocking, swarming, splitting, and cooperative coverage) without exposing expensive hardware to danger.

Therefore, DT frameworks significantly expand the design space for UAV autonomy. When combined with SITL, they support full-stack testing: from perception and control to networking and inter-agent coordination, forming the technological foundation for this paper's proposed DT-SITL architecture.

C. Software-in-the-Loop Simulation Approaches

Software-in-the-Loop (SITL) simulation enables developers to execute flight control software, autonomy modules, or AI decision systems in a fully virtual environment without requiring access to physical UAV hardware. SITL bridges the gap between purely algorithmic simulations and hardware deployment by incorporating actual firmware or platform-specific control software into the simulation loop.

Common SITL frameworks include PX4 SITL, ArduPilot SITL, and custom SITL pipelines integrated with simulators like AirSim or Gazebo. As explained in Dimmig et al. (2025), SITL systems can be configured to simulate noisy sensor inputs, actuator delays, environmental disturbances, and realistic vehicle dynamics. This enables robust testing of:

- navigation control loops
- autonomous mission execution
- failure handling and recovery
- swarm management and distributed planning
- dynamic object tracking and avoidance

The use of SITL is especially crucial for multi-agent UAV systems, where subtleties in timing, network latency, or sensor noise can significantly impact emergent behaviors. For surveillance, SITL allows repeated testing of coordinated maneuvers (such as cooperative pursuit, perimeter coverage, or synchronized waypoint following) under identical virtual conditions.

Recent platforms such as U2USim (Han et al., 2024) demonstrate how SITL can be extended to include telepresence, high-fidelity sensing, and dynamic environmental modeling. Similarly, integrated DT-SITL systems have been used to evaluate autonomous flight policies in complex environments, such as landing in turbulent conditions (McClellan et al., 2021) or swarm-wide communication resilience (Zhou et al., 2024).

In digital twin workflows, SITL often acts as the execution layer, where the digital model interacts with real-world data streams and control commands. This provides a complete loop for validating algorithms end-to-end, including perception, control, navigation, and networking.

Overall, SITL provides an indispensable foundation for UAV autonomy research. Its repeatability, safety, and low cost make it integral to this study's proposed architecture, enabling multi-agent DT-driven UAV simulation with realistic physics and AI-in-the-loop coordination.

While these areas (multi-agent UAV coordination, digital twin architectures, and software-in-the-loop) have matured independently, several limitations remain evident. First, most digital twin implementations are tailored to single-UAV scenarios, with limited support for coordinated sensing, cooperative decision-making, or distributed multi-drone task execution. Second, existing simulators often prioritize either photorealistic perception or flight-controller fidelity, but comparatively few provide an integrated environment capable of supporting perception-rich, multi-agent intelligence. Third, current workflows seldom combine DT, SITL, and autonomous coordination within a unified simulation-to-deployment pipeline.

These gaps collectively underscore the need for a framework that brings together high-fidelity perception, synchronized digital twinning, and multi-agent autonomy into a cohesive and scalable system. Motivated by these findings, Section III introduces the proposed Digital Twin–Software-in-the-Loop (DT-SITL) framework, designed to support multi-UAV surveillance missions under coordinated, distributed intelligence. The framework begins with a comparative analysis of state-of-the-art UAV simulation tools, establishing the rationale for selecting AirSim as the primary environment. It then expands into the architectural design, synchronization mechanisms, and multi-agent modules that enable fully coordinated aerial surveillance in a DT-driven SITL ecosystem.

III. Digital Twin–SITL Framework and Learning-Based Multi-Agent Architecture

This section presents the proposed Digital Twin–Software-in-the-Loop (DT–SITL) framework for coordinated multi-UAV surveillance and tracking. The framework integrates digital twin modeling, learning-based multi-agent control, and synchronized simulation execution to enable safe, repeatable, and scalable validation of cooperative UAV behaviors prior to physical deployment.

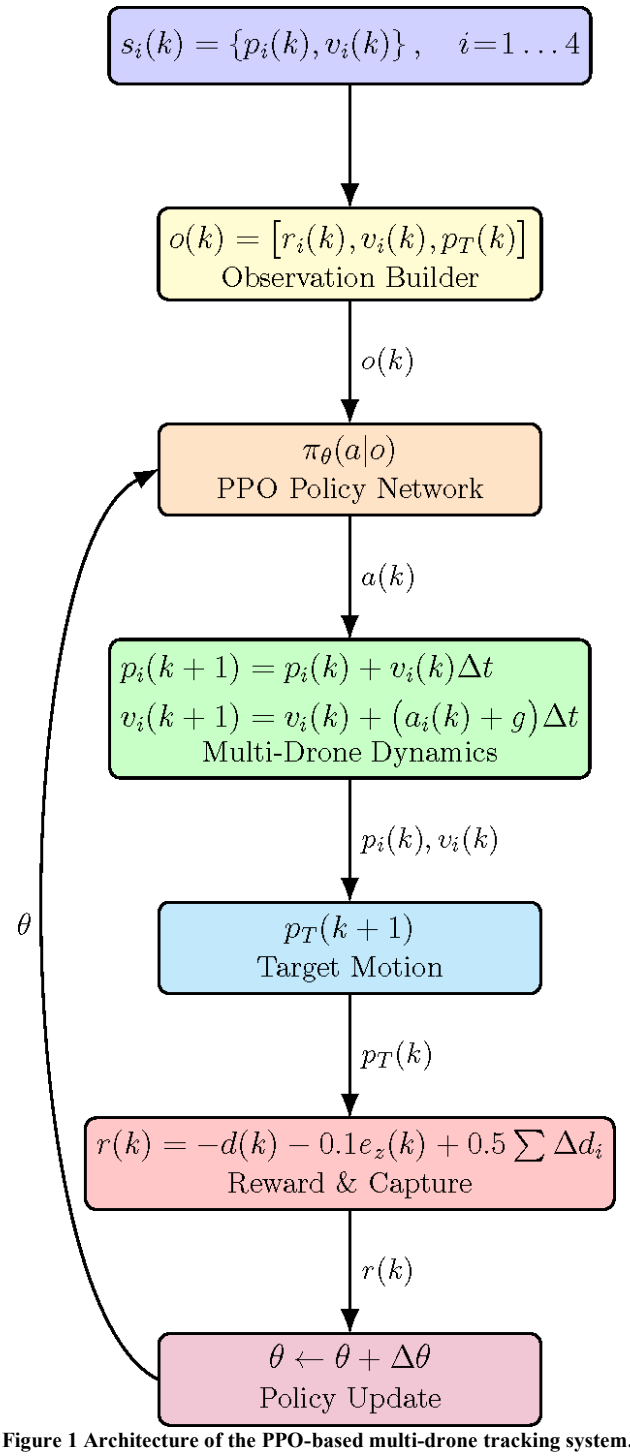

Figure 1 Architecture of the PPO-based multi-drone tracking system.

Unlike conventional UAV simulation pipelines that treat vehicle dynamics, autonomy logic, and mission evaluation as loosely coupled components, the proposed DT–SITL framework embeds learning, perception abstraction, and coordination incentives directly into the digital twin loop. Each UAV is associated with a dedicated DT instance and controlled by a shared PPO-based policy, enabling decentralized execution while preserving coordinated system-level behavior.

Figure 1 illustrates the overall architecture of the PPO-based multi-drone tracking system within the DT–SITL framework.

A. UAV Dynamics and Digital Twin State Representation

Each UAV is represented by a DT instance whose state evolves in discrete time. The state of UAV i at time step k is defined as:

$$s_i(k) = \begin{bmatrix} p_i(k) \\ v_i(k) \end{bmatrix} \in \mathbb{R}^6 \quad \textbf{(1)}$$

where $p_i(k) = [x_i, y_i, z_i]^{\mathrm{T}}$ denotes position and $v_i(k) = [v_{x\,i}, v_{y\,i}, v_{z\,i}]^{\mathrm{T}}$ denotes velocity.

The DT dynamics follow a discrete-time double-integrator model:

$$v_i(k+1) = v_i(k) + (\, a_i(k) + g \,) \cdot \Delta t \quad \textbf{(2)}$$

$$p_i(k+1) = p_i(k) + v_i(k) \cdot \Delta t \quad \textbf{(3)}$$

where $a_i(k)$ is the acceleration command generated by the agent policy and g represents gravitational acceleration. This abstraction balances computational efficiency with sufficient fidelity for coordinated tracking behavior.

B. Target Motion Model and Environment Interaction

The DT–SITL environment includes a dynamically maneuvering target whose trajectory is explicitly modeled to induce cooperative behavior among UAVs. The target position is defined as a smooth, time-varying function in three-dimensional space:

$$p_T(k) = \begin{bmatrix} 40 \cdot \sin(0.05 \cdot k \cdot \Delta t) \\ 40 \cdot \cos(0.04 \cdot k \cdot \Delta t) \\ 20 + 8 \cdot \sin(0.03 \cdot k \cdot \Delta t) \end{bmatrix} \quad \textbf{(4)}$$

This motion profile introduces non-uniform acceleration and altitude variation, forcing UAV agents to continuously adapt their trajectories and maintain spatial coordination. By embedding target dynamics within the DT layer, the framework enables closed-loop agent–environment interaction suitable for learning-based surveillance tasks.

C. Observation Construction and Perception Abstraction

Rather than exposing raw sensor data to the learning agent, the DT–SITL framework constructs a structured observation vector that captures relative spatial relationships critical for cooperative tracking.

For UAV i, the relative displacement to the target is defined as:

$$r_i(k) = p_i(k) - p_T(k) \quad \textbf{(5)}$$

The global observation provided to the PPO agent is:

$$o(k) = \begin{bmatrix} r_1(k) \\ v_1(k) \\ r_2(k) \\ v_2(k) \\ r_3(k) \\ v_3(k) \\ r_4(k) \\ v_4(k) \\ p_T(k) \end{bmatrix} \in \mathbb{R}^{27} \quad \textbf{(6)}$$

This abstraction reflects a digital twin–level perception layer, decoupling sensing fidelity from decision-making while preserving the spatial coupling required for coordination.

D. Action Space, Policy Structure, and Reward Design

The agent outputs acceleration commands for all UAVs:

$$a(k) = \begin{bmatrix} a_1(k) \\ a_2(k) \\ a_3(k) \\ a_4(k) \end{bmatrix} \in \mathbb{R}^{12} \quad \textbf{(7)}$$

where each drone receives:

$$a_i(k) = \begin{bmatrix} a_{x\,i}(k) \\ a_{y\,i}(k) \\ a_{z\,i}(k) \end{bmatrix} \in \mathbb{R}^3 \quad \textbf{(8)}$$

Accelerations are saturated element-wise:

$$a_i(k) = \mathrm{clip}\big(\, \hat{a}_i(k),\ -a_{\max},\ a_{\max} \,\big) \quad \textbf{(9)}$$

The PPO policy is modeled as a Gaussian distribution:

$$\pi_\theta(a|o) = \mathcal{N}(\, W_\pi^{\mathrm{T}} \cdot o,\ \Sigma \,) \quad \textbf{(10)}$$

The reward function explicitly encodes cooperative tracking objectives:

$$r(k+1) = -d(k+1) - 0.1 \cdot e_z(k+1) + 0.5 \cdot \sum_i \Delta d_i(k) \quad \textbf{(11)}$$

where $d(k)$ denotes mean UAV–target distance and $e_z(k)$ penalizes altitude mismatch. This formulation discourages independent pursuit and promotes coordinated spatial distribution.

E. DT–SITL Execution Loop and Metric-Aware Design

Figure 2 illustrates the conceptual SITL architecture integrating policy execution, dynamics propagation, reward computation, and logging.

To ensure deterministic execution and avoid overlapping abstractions, the implemented system adopts the streamlined architecture shown in Figure 3.

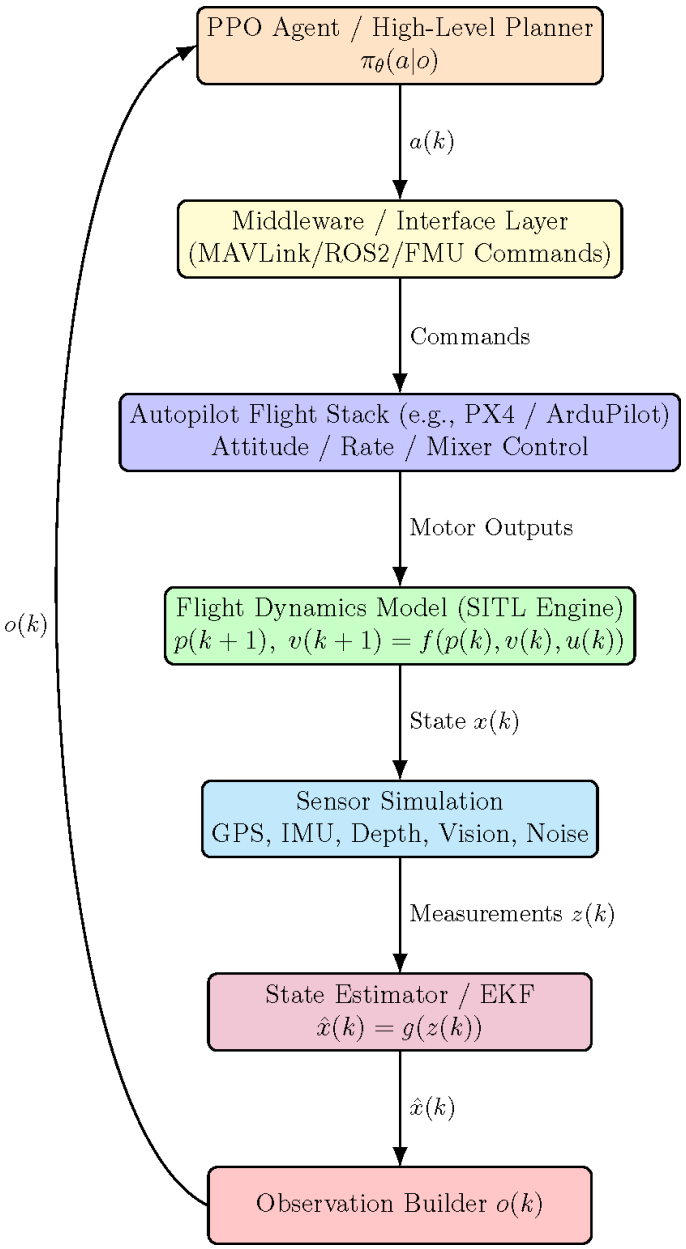

Figure 2 Conceptual SITL architecture for the multi-drone PPO tracking system.

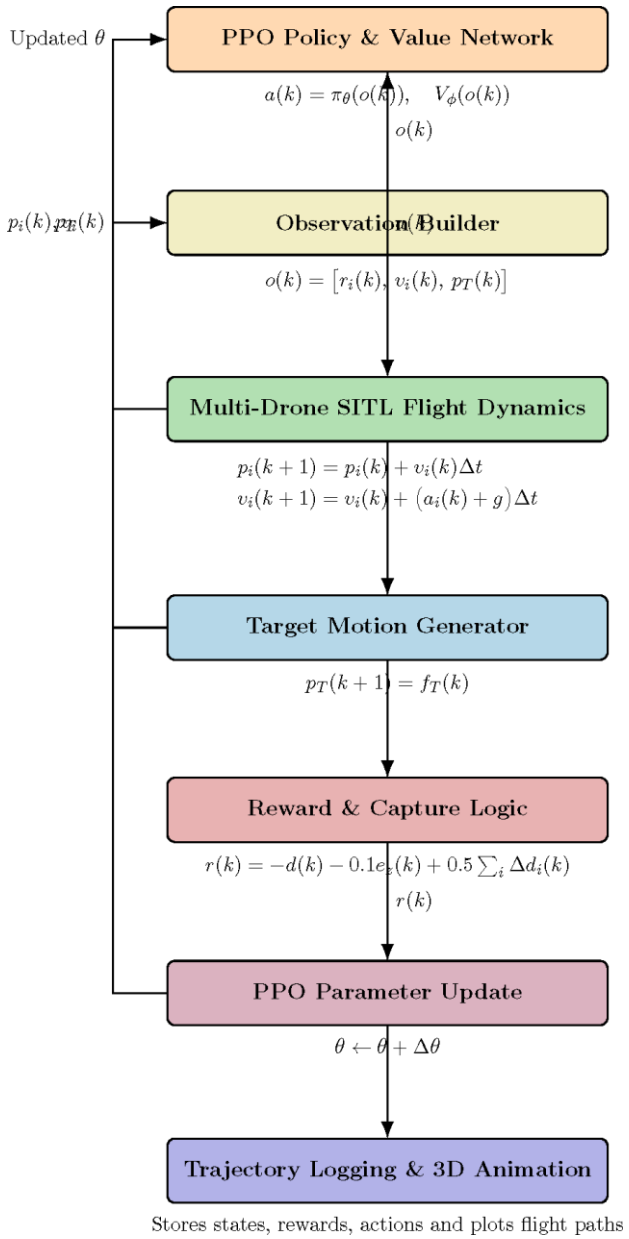

Figure 3 Implemented SITL architecture with non-overlapping mathematical models.

Importantly, the DT–SITL loop is explicitly designed to expose evaluation-relevant signals, including:

- UAV and target trajectories
- Position error over time
- Inter-agent distances
- Policy reward evolution

These signals directly inform the metrics analyzed in Section IV, ensuring that experimental results are not post-hoc artifacts but intrinsic outputs of the DT–SITL execution.

IV. Results and Discussion

This section presents the results obtained from the proposed DT–SITL framework and discusses their implications for coordinated multi-UAV surveillance and tracking. The evaluation focuses on emergent multi-agent behavior, including tracking accuracy, coordination effectiveness, and learning convergence, as observed during simulation-based missions involving three UAVs.

A. Mission Geometry and Surveillance Trajectory

The UAVs follow a predefined rectangular surveillance trajectory, establishing a structured mission environment for evaluating cooperative tracking behavior.

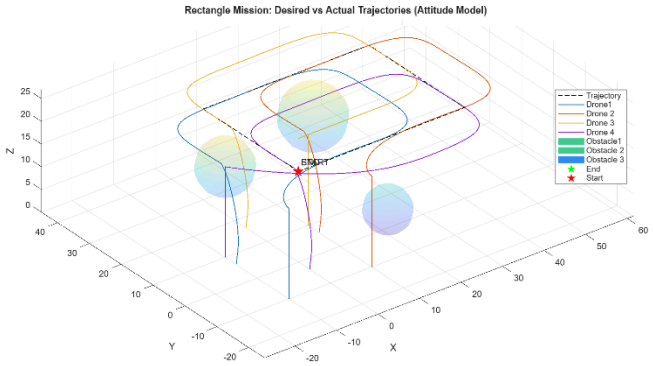

Figure 4 Rectangular surveillance trajectory executed by the multi-UAV system.

B. *Tracking Accuracy and Position Error*

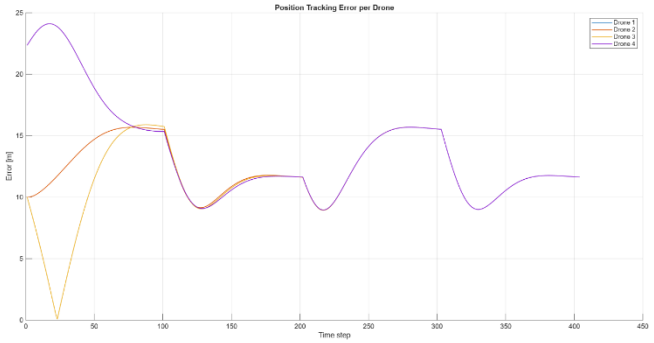

Figure 5 Position error during coordinated tracking.

The bounded position error across UAVs indicates stable closed-loop behavior and effective compensation for dynamic target motion.

C. *Learning Convergence and Reward Evolution*

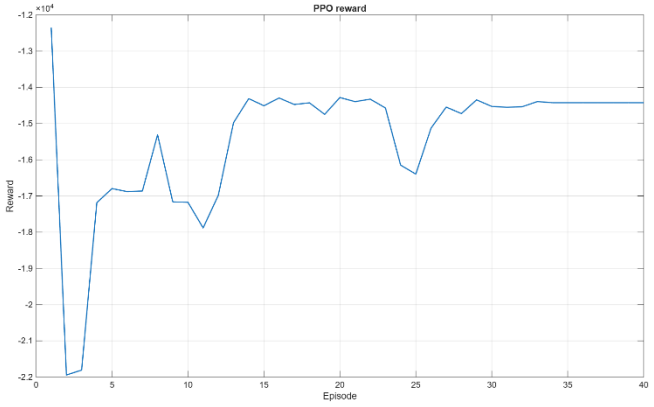

Figure 6 PPO reward evolution during training.

The increasing reward trend demonstrates successful policy learning and convergence toward coordinated tracking strategies.

D. *Distance Regulation and Cooperative Engagement*

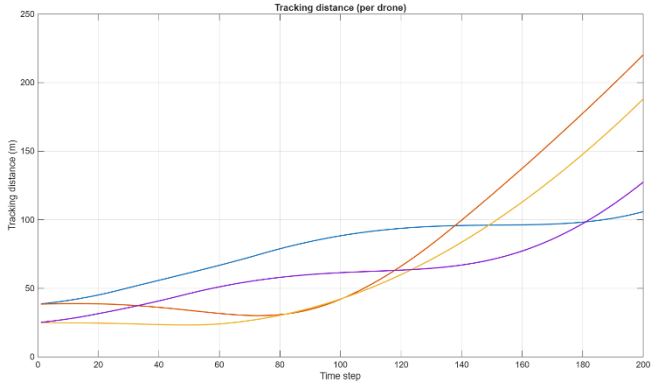

Figure 7 Mean tracking distance between UAVs and the target.

Distance regulation reflects coordinated spatial distribution rather than independent pursuit.

E. *Emergent Multi-UAV Tracking Behavior*

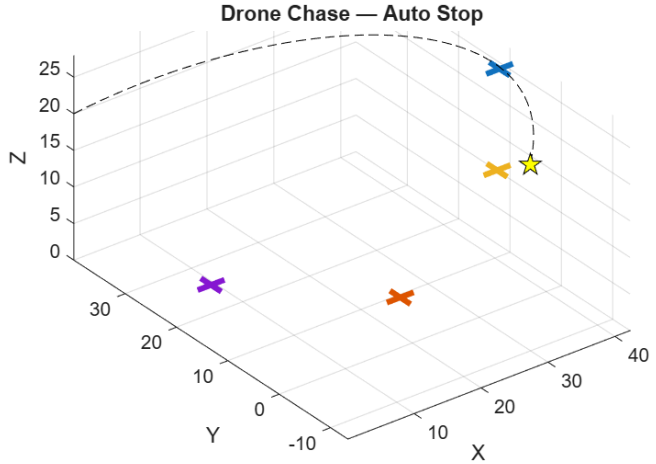

Figure 8 Moving target tracking performance.

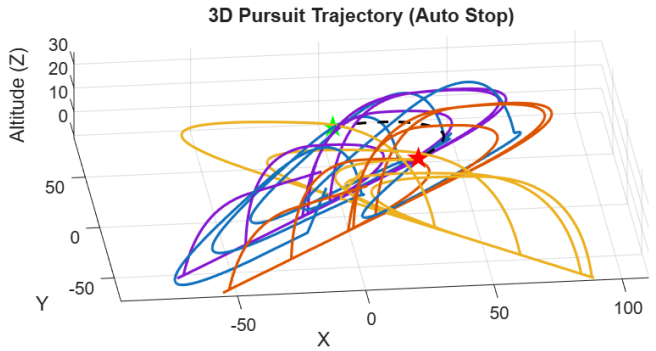

Figure 9 Coordinated multi-UAV tracking trajectories.

These results highlight emergent cooperation, persistent engagement, and smooth coordination induced by the shared learning-based policy.

The results presented above demonstrate stable learning, accurate tracking, and effective coordination under dynamic mission conditions. Beyond individual performance metrics, these results reveal system-level properties that emerge from the interaction between Digital Twin modeling, learning-based control, and decentralized execution. These properties are analyzed in the following section.

V. System-Level Analysis and Implications

Building on the results and observations presented in Section IV, this section analyzes the system-level behavior of the proposed DT–SITL framework and discusses its implications for scalable and autonomous aerial surveillance systems.

At the system level, three key properties emerge:

1. Digital Twin–mediated consistency, which stabilizes learning and coordination
2. Emergent robustness, arising from multi-agent redundancy
3. Scalability potential, enabled by decentralized execution with shared policy structure

The DT–SITL framework transforms simulation from a passive testing tool into an active co-evolving component of system design, enabling policy validation, coordination analysis, and safety assessment within a unified loop.

These findings suggest that DT–SITL–based learning architectures provide a viable foundation for future large-scale, autonomous, and resilient UAV surveillance deployments.

VI. Potential Applications

While the proposed DT–SITL framework is evaluated in a controlled multi-UAV tracking scenario, its architectural intent is broader: it provides a repeatable simulation-to-deployment pipeline for developing and validating cooperative aerial behaviors under safety, scalability, and coordination constraints. This subsection outlines representative real-world application domains where a DT–SITL approach with learning-based multi-agent control is particularly relevant, and clarifies the operational value of the proposed design.

A. Urban security and public-safety surveillance.

Dense urban environments present occlusions, constrained flight corridors, and rapidly changing target motion patterns. Multi-UAV teams can maintain persistent situational awareness by distributing vantage points, reducing blind spots and improving continuity of tracking. In this context, the DT layer enables scenario replay (e.g., recurring crowd flows, road-network constraints, and typical target maneuvers), while SITL supports end-to-end evaluation of autonomy logic without flight risk. The shared-policy cooperative tracking developed in this work is applicable to tasks such as suspect tracking across intersections, monitoring high-density public events, and maintaining coverage of critical urban perimeters. Importantly, DT–SITL enables systematic stress testing (e.g., occlusion spikes, route bottlenecks) before operational deployment, improving safety assurance and mission predictability.

B. Search-and-rescue (SAR) and disaster response.

SAR missions demand rapid coverage and resilient coordination under uncertain conditions, where communication may degrade and terrain may be partially unknown. Cooperative behaviors such as persistent engagement, adaptive spacing, and distributed pursuit are useful for tracking a moving subject, guiding rescuers, or coordinating multi-view confirmation. The DT–SITL framework is well-suited to SAR training because it enables controlled injection of disturbances (e.g., sensor noise, wind, intermittent links) and repeatable evaluation across standardized scenarios. Moreover, DT-driven scenario rehearsal allows agencies to validate mission logic (coverage strategies, handover protocols, and cooperative pursuit) in a risk-free environment prior to field execution, thereby reducing time-to-deploy and improving operational robustness.

C. Infrastructure inspection and industrial asset monitoring.

Industrial sites (pipelines, ports, large plants, rail corridors, and bridges) require persistent inspection with strict safety constraints and complex geometry. Multi-UAV systems can coordinate to maintain standoff distances, share line-of-sight responsibilities, and provide continuous monitoring of moving assets (e.g., vehicles, vessels, or mobile machinery). DT–SITL is particularly beneficial here because a digital twin of the asset environment (layout, no-fly zones, and inspection waypoints) can be used for mission rehearsal and for evaluating cooperative behaviors such as dynamic target following and distributed coverage. The proposed framework supports iterative refinement of reward structures and observation abstractions to match domain objectives (e.g., maintain camera geometry, maximize coverage continuity), while preserving repeatability for performance comparison across mission configurations.

D. Border, perimeter, and critical-facility protection.

Perimeter protection often requires persistent tracking of intrusions across large areas with heterogeneous terrain and intermittent visibility. Multi-UAV coordination improves resilience through redundancy: if one UAV loses line-of-sight, others can maintain engagement. The DT–SITL architecture enables agencies to test cooperative pursuit and handover behaviors, assess coordination stability, and evaluate response policies prior to deployment. In addition, DT–SITL supports scalable design-space exploration—such as varying team sizes, target behaviors, and patrol geometries—allowing stakeholders to identify operational limits and safety margins without risking costly assets.

E. Scalable coordination analysis and transition toward deployment.

Across these applications, the key practical benefit of DT–SITL is the ability to validate multi-agent autonomy at the system level—tracking performance, coordination stability, and emergent behaviors—under controlled and repeatable conditions. The proposed framework provides a natural

pathway for progressive realism: starting from structured observation abstractions for stable learning and scaling, then incorporating higher-fidelity sensing, disturbances, and communication constraints as the system transitions toward hardware-in-the-loop (HITL) and field testing. Consequently, DT–SITL functions not only as a simulator but as an engineering workflow that bridges algorithm development, verification, and operational readiness for multi-UAV surveillance systems.

VII. Conclusion

This paper presented a Digital Twin–Software-in-the-Loop (DT–SITL) framework for coordinated multi-UAV aerial surveillance, integrating learning-based multi-agent control with synchronized simulation execution. By embedding a shared Proximal Policy Optimization (PPO) agent within a Digital Twin–driven SITL architecture, the proposed approach enables systematic development, validation, and analysis of cooperative tracking behaviors prior to physical deployment.

Unlike conventional UAV simulation pipelines that decouple dynamics, autonomy, and evaluation, the proposed framework tightly integrates these components into a unified execution loop. Each UAV is represented by a dedicated Digital Twin instance, while decision-making is governed by a shared learning-based policy operating on structured observations derived from relative spatial relationships. This design enables decentralized execution with emergent coordinated behavior, preserving scalability while maintaining system-level consistency.

Simulation results demonstrated that the proposed DT–SITL framework achieves stable learning convergence, bounded tracking error, and effective multi-UAV coordination when tracking a dynamically maneuvering target. The results further showed that cooperative behaviors (such as spatial distribution, distance regulation, and persistent engagement) emerge naturally from the reward structure without explicit formation constraints. These findings confirm that the integration of Digital Twin modeling with learning-based control is a viable and effective strategy for multi-agent aerial surveillance.

From a system-level perspective, the DT–SITL framework offers several important advantages. First, it enables safe and repeatable experimentation, allowing complex multi-agent behaviors to be evaluated under controlled conditions without risking physical assets. Second, it supports scalable design-space exploration, making it possible to assess coordination strategies, reward formulations, and mission geometries before deployment. Third, it provides a foundation for progressive autonomy, facilitating future transitions from SITL to hardware-in-the-loop (HITL) and real-world multi-UAV operations.

Despite these strengths, several limitations remain. The current study focuses on simulation-based evaluation with simplified vehicle dynamics and a limited number of UAVs. Environmental disturbances such as wind, sensor noise, and communication interference are abstracted rather than fully modeled. In addition, the shared-policy assumption, while effective for coordination, may limit heterogeneity across UAV roles in more complex missions.

Future work will address these limitations by extending the framework to larger UAV teams, incorporating higher-fidelity aerodynamic and sensor models, and integrating communication-aware learning mechanisms. The DT–SITL framework also provides a natural pathway toward federated or decentralized learning, heterogeneous agent roles, and real-world validation using hardware-in-the-loop and flight testing. Overall, this work establishes DT–SITL–enabled learning architectures as a promising foundation for next-generation autonomous aerial surveillance systems.

References

[1] Ahmadzadeh, Ali, Ali Jadbabaie, Vijay Kumar, and George J. Pappas. “Multi-UAV Cooperative Surveillance with Spatio-Temporal Specifications”. In *Proceedings of the 45th IEEE Conference on Decision and Control*, 5293–98, 2006. https://doi.org/10.1109/CDC.2006.377157.

[2] Aláez, D., P. Lopez-Iturri, M. Celaya-Echarri, L. Azpilicueta, F. Falcone, J. Villadangos, and J. J. Astrain. “Digital Twin Modeling of Open Category UAV Radio Communications: A Case Study”. *Computer Networks* 242 (2024): 110276. https://doi.org/10.1016/j.comnet.2024.110276.

[3] Danish Iqbal and Barbora Buhnova, “Digital Twin Design for Autonomous Drones”, in *Proceedings of the 19th Conference on Computer Science and Intelligence Systems* (FedCSIS). *Annals of Computer Science and Information Systems* vol.39 (2024), pp. 119–130. https://doi.org/10.15439/2024F6765.

[4] Dimmig, Cora A., Giuseppe Silano, Kimberly McGuire, Chiara Gabellieri, Wolfgang Hönig, Joseph Moore, and Marin Kobilarov. “Survey of Simulators for Aerial Robots: An Overview and In-Depth Systematic Comparisons”. *IEEE Robotics & Automation Magazine* 32, no. 2 (2025): 153–66. https://doi.org/10.1109/MRA.2024.3433171.

[5] Ekechi, Chijioke C., Tarek Elfouly, Ali Alouani, and Tamer Khattab. “A Survey on UAV Control with Multi-Agent Reinforcement Learning”. *Drones* 9, no. 7 (2025). https://doi.org/10.3390/drones9070484.

[6] Gray, William, II, Abduallah Al Redwan Newaz, Milad Khaleghi, Hieu T. Nguyen, Leila Hashemi Beni, Abolghasem Shahbazi, and Ali Karimoddini. “Multi-UAV Tasking and Coordination for Monitoring and Coverage of Agricultural Farmland”. *ACM J. Auton. Transport. Syst.* 3, no. 1 (September 2025). https://doi.org/10.1145/3749194.

[7] Gürses, Anıl, Gautham Reddy, Saad Masrur, Özgür Özdemir, İsmail Güvenç, Mihail L. Sichitiu, Alphan Şahin, Ahmed Alkhateeb, Magreth Mushi, and Rudra Dutta. “Digital Twins and Testbeds for Supporting AI Research with Autonomous Vehicle Networks”. *arXiv [Eess.SP]*, 2024. arXiv. http://arxiv.org/abs/2404.00954.

[8] Han, Feilin, Leping Zhang, Xin Wang, Ke-Ao Zhao, Ying Zhong, Ziyi Su, Tongtong Feng, and Wenwu Zhu. “U2USim - A UAV Telepresence Simulation Platform with Multi-Agent Sensing and Dynamic Environment”. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 11258–60. MM ’24. New York, NY, USA: Association for Computing Machinery, 2024. https://doi.org/10.1145/3664647.3684995.

[9] Liang, Zibin, Qing Li, and Guodong Fu. “Multi-UAV Collaborative Search and Attack Mission Decision-Making in Unknown Environments”. *Sensors (Basel, Switzerland)* 23, no. 17 (24 August 2023). https://doi.org/10.3390/s23177398.

[10] McClellan, Andrew R., Joseph Lorenzetti, Marco Pavone, and Charbel Farhat. 2021. *A Physics-Based Digital Twin for Model Predictive Control of Autonomous Unmanned Aerial Vehicle Landing.* Stanford Autonomous Systems Lab. https://stanfordasl.github.io/wp-content/papercite-data/pdf/McClellan.Lorenzetti.ea.PTRSA21.pdf.

[11] Seraj, Esmaeil, Andrew Silva, and Matthew Gombolay. “Multi-UAV Planning for Cooperative Wildfire Coverage and Tracking with Quality-of-Service Guarantees”. *arXiv [Cs.RO]*, 2022. arXiv. http://arxiv.org/abs/2206.10544.

[12] Su, Kai, and Feng Qian. “Multi-UAV Cooperative Searching and Tracking for Moving Targets Based on Multi-Agent Reinforcement Learning”. *Applied Sciences* 13, no. 21 (2023). https://doi.org/10.3390/app132111905.

[13] Wang, Yi’an, Kun Li, Ying Han, and Xinxin Yan. “Distributed Multi-UAV Cooperation for Dynamic Target Tracking Optimized by an SAQPSO Algorithm”. *ISA Transactions* 129 (2022): 230–42. https://doi.org/10.1016/j.isatra.2021.12.014.

[14] Yahya, Muhammad, Muhammad Naeem, Zeeshan Kaleem, and Waleed Ejaz. “Robust Priority Aware Multi-Criterion Offloading in Digital Twin UAVs Networks”. *Internet of Things* 34 (2025): 101763. https://doi.org/10.1016/j.iot.2025.101763.

[15] Zami, Md Bokhtiar Al, Md Raihan Uddin, and Dinh C. Nguyen. “Secure UAV-Assisted Federated Learning: A Digital Twin-Driven Approach with Zero-Knowledge Proofs”. *arXiv [Cs.LG]*, 2025. arXiv. http://arxiv.org/abs/2509.13634.

[16] Zhou, Longyu, Supeng Leng, and Qing Wang. “A Federated Digital Twin Framework for UAVs-Based Mobile Scenarios”. *IEEE Transactions on Mobile Computing* 23, no. 6 (2024): 7377–93. https://doi.org/10.1109/TMC.2023.3335386.