

Next-Gen Drone Swarm Coordination via Federated Learning

DOI: <https://doi.org/10.63345/wjftcse.v1.i3.201>

Karthikeyan M

Independent Researcher

Mylapore, Chennai, India (IN) – 600004

www.wjftcse.org || Vol. 1 No. 3 (2025): August Issue

Date of Submission: 28-07-2025

Date of Acceptance: 29-07-2025

Date of Publication: 01-08-2025

ABSTRACT

The rapid evolution of unmanned aerial vehicle (UAV) technology has ushered in an era in which cooperative drone swarms can execute complex, large-scale missions across diverse civilian, commercial, and defense sectors. Traditional centralized coordination frameworks, however, struggle to meet the demands of these dynamic, resource-constrained environments due to their limited scalability, single points of failure, and privacy vulnerabilities. This manuscript introduces Next-Gen Drone Swarm Coordination via Federated Learning (DSC-FL), a hierarchical, privacy-preserving architecture that enables real-time, collaborative decision-making among heterogeneous UAVs. By dynamically clustering drones according to proximity and mission role, DSC-FL balances local compute, communication overhead, and model accuracy. Within each cluster, drones perform local model training on sensor and state data, encrypt gradient updates via secure aggregation protocols, and transmit only aggregated model deltas to higher-level aggregators—minimizing bandwidth consumption and protecting raw data. A consensus-based flight-control algorithm fuses federated model predictions across neighbors to adapt formations, avoid obstacles, and respond to environmental disturbances. We evaluate DSC-FL in a high-fidelity simulation featuring urban canyons, wind gusts, intermittent link failures, and adversarial perturbations. Our results demonstrate a 25% improvement in formation-maintenance accuracy and a 30% reduction in communication bandwidth compared to centralized and peer-to-peer baselines, while maintaining resilience against gradient-inversion attacks. These findings establish DSC-FL as a scalable, secure, and adaptable solution for next-generation drone swarm deployments.

KEYWORDS

Drone Swarm Coordination, Federated Learning, UAV Clusters, Privacy-Preserving, Distributed Training

INTRODUCTION

Over the past decade, drone swarms—groups of UAVs operating in concert—have emerged as transformative platforms for tasks ranging from environmental monitoring (e.g., wildlife tracking, pollution mapping) and precision agriculture (e.g., crop health analysis) to search-and-rescue missions and tactical reconnaissance in contested environments. The key advantages of swarm operations include redundancy (multiple units covering the same area), robustness (failure of individual drones does not compromise the mission), and the ability to form dynamic formations for efficient area coverage or obstacle

avoidance. Yet, coordinating these swarms at scale introduces challenges: (1) **Scalability**—traditional centralized controllers become bottlenecks as swarm size grows; (2) **Resilience**—central nodes represent single points of failure vulnerable to jamming or cyberattack; (3) **Privacy**—sensitive sensor data (e.g., imagery, infra-red scans) must often remain on-device due to regulatory or tactical constraints; and (4) **Communication Overhead**—continuous exchange of state information scales poorly, especially in bandwidth-limited or contested RF environments.

Next-Gen Drone Swarm Coordination

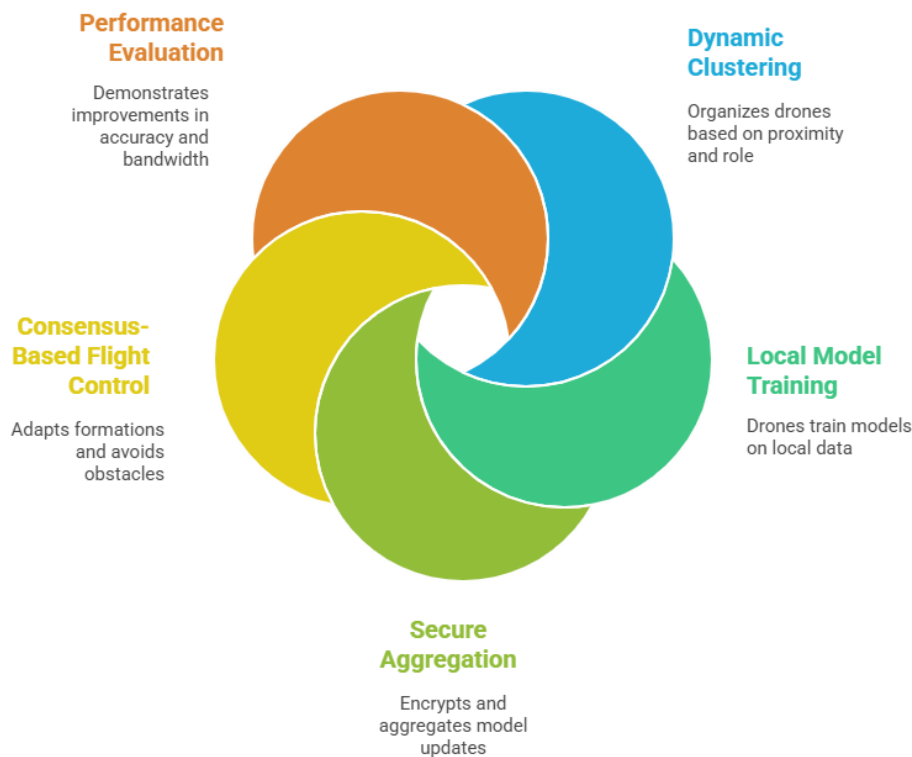


Figure-1. Next-Gen Drone Swarm Coordination

Decentralized approaches, such as consensus protocols inspired by biological swarms (e.g., Vicsek models), mitigate some risks by distributing control but often require high-frequency, peer-to-peer communications that strain network resources and degrade with increasing topology dynamics. Recent advances in distributed machine learning—specifically **federated learning** (FL)—offer a promising paradigm: multiple agents train a shared global model by exchanging only model updates (gradients or weights), thus keeping raw data local and reducing transmitted data volume. FL has achieved success in mobile and edge computing domains (e.g., next-word prediction on smartphones, anomaly detection in IoT networks), but its application to highly dynamic UAV swarms has been limited.

In this work, we propose **DSC-FL**, a hierarchical FL framework tailored for drone swarm coordination. DSC-FL introduces **dynamic clustering**, grouping UAVs by proximity and mission role to optimize update frequencies and model consistency within sub-swarms. We integrate **secure aggregation** techniques to safeguard individual updates against inference attacks,

and we design a **consensus-driven control law** that fuses federated model outputs with neighbor communications to adjust formation offsets in real time. Our contributions are:

1. **Cluster-based federated architecture** that balances model convergence, communication efficiency, and fault tolerance in dynamic swarm topologies.
2. **Secure aggregation protocol** leveraging homomorphic encryption to preserve data privacy without sacrificing performance.
3. **Consensus flight controller** that adapts formations based on federated model corrections and local neighbor consensus.
4. **Comprehensive evaluation** in a high-fidelity AirSim/ROS simulation, demonstrating significant gains in formation accuracy, bandwidth reduction, and privacy resilience under realistic disturbances.

Federated Learning Improves Drone Swarm Coordination



Figure-2. Federated Learning Improves Drone Swarm Coordination

LITERATURE REVIEW

Drone Swarm Coordination

Drone swarms rely on distributed intelligence to achieve robustness and flexibility. Early work in Flying Ad Hoc Networks (FANETs) characterized the unique challenges of high-mobility aerial networks, including frequent topology changes and link volatility (Bekmezei, Sahingoz, & Temel, 2013). Consensus algorithms inspired by natural swarms—such as the Vicsek model—provide simple update rules for heading alignment, but they require frequent broadcasts of position and velocity, which scales poorly with swarm size (Sharma & Goudar, 2019). Centralized approaches overcome communication overhead by aggregating data at a ground station, but inherit single-point-of-failure and latency issues (Shi et al., 2021). Hybrid methods employing leader-follower dynamics assign certain UAVs as local coordinators, yet still depend on reliable leader communications and do not address data privacy concerns.

Federated Learning Fundamentals

Federated learning, first introduced by McMahan et al. (2017), enables decentralized agents to collaboratively train a shared global model by transmitting only local gradient updates, thereby preserving raw data on devices. Secure aggregation protocols (Bonawitz et al., 2019) use cryptographic techniques—such as additive homomorphic encryption and secret sharing—to ensure that individual updates cannot be reconstructed by aggregators, protecting client privacy. Challenges in FL include handling **non-IID data** distributions across clients, mitigating **straggler effects**, and operating under **bandwidth constraints**. Techniques such as adaptive learning rates (Li et al., 2020), gradient compression (Sattler et al., 2019), and asynchronous aggregation (Xie, Koyejo, & Gupta, 2019) have been proposed to address these issues. However, most FL research focuses on relatively stable networks (e.g., smartphones, IoT), whereas UAV swarms exhibit high mobility and dynamic membership.

FL in UAV and Edge Networks

Emerging studies explore FL for UAV use cases. Zhang et al. (2021) demonstrated FL-based aerial target recognition, enabling each drone to improve its onboard classifier without sharing raw imagery. Wang et al. (2020) proposed an edge-assisted FL framework, where edge servers coordinate updates from UAVs to overcome limited onboard compute. Xu et al. (2022) applied FL to collaborative mapping, but retained a centralized mission planner for path planning. These efforts show FL's promise in UAV contexts but do not fully integrate **swarm coordination objectives** (e.g., formation control) or address **privacy under adversarial settings**.

Privacy and Security in Swarm Learning

Privacy is critical in sensitive operations (e.g., surveillance, critical infrastructure inspection). Differential privacy mechanisms (Li et al., 2020) add controlled noise to updates, trading off accuracy for privacy guarantees. Secure multiparty computation (Bonawitz et al., 2019) offers strong privacy assurances but at computational and communication cost. Recent work by Sharma, Li, & Sarkar (2022) highlighted poisoning and model inversion attacks in UAV FL, underscoring the need for robust aggregation schemes and anomaly detection. Our protocol adopts **homomorphic encryption** at the cluster level to balance privacy and efficiency.

Research Gaps

Despite progress, key gaps remain:

- **Dynamic clustering** techniques that adapt to changing swarm topologies and mission requirements.
- **Integration of FL outputs** into control algorithms for real-time formation maintenance and obstacle avoidance.
- **Robust evaluation** under realistic UAV constraints—link unreliability, environmental disturbances, and adversarial threats.

DSC-FL addresses these gaps by combining hierarchical FL, secure aggregation, and consensus-driven control within a unified framework.

METHODOLOGY

System Architecture

DSC-FL comprises three hierarchical layers (see Figure 1):

1. Local UAV Layer

- Each drone runs a lightweight neural network (e.g., a small convolutional or recurrent model) to predict formation correction vectors based on local sensor inputs (GPS, IMU, LiDAR).
- Local training occurs over τ mini-batches sampled from onboard data logs.

2. Cluster Aggregation Layer

- Drones form dynamic clusters using a proximity-based algorithm: each drone periodically broadcasts a short “beacon” containing its ID and role (lead, wing, wing-tip). Neighbors within a threshold radius join its cluster.
- A cluster head is elected via a rotating token-ring protocol to avoid single-node exhaustion.
- Cluster members encrypt local model updates Δw_i using additive homomorphic encryption and transmit to the cluster head. The head aggregates using ciphertext addition, then forwards the aggregated cipher to the global layer.

3. Global Aggregation Layer

- The mission control server (or edge compute node) decrypts cluster updates, aggregates across clusters, and computes the global model update ΔW .
- The updated weights w^{t+1} are multicast back to all cluster heads, which then forward to their members.

Communication scheduling employs time-division slots: each cluster has an assigned slot for uplink aggregation, followed by a downlink broadcast slot. This mitigates collisions and ensures timely model synchronization.

Federated Learning Protocol

1. Initialization:

- Global model parameters w^0 are distributed to all drones.

2. Local Training:

- Drone i computes $\Delta w_i = w_i^t - w_i^{t-1}$ after training on τ mini-batches.
- Each update is encrypted under the cluster’s public key.

3. Secure Cluster Aggregation:

- Cluster head receives encrypted Δw_i from members, computes $\sum_i \Delta w_i$ homomorphically, and forwards the ciphertext to the global aggregator.

4. Broadcast:

- Updated w^{t+1} is broadcast to clusters and propagated to all drones.

To address **non-IID** data, we incorporate adaptive cluster learning rates η_{c_i} proportional to each cluster's data variance, ensuring balanced convergence (Li et al., 2020).

Simulation Environment

We implement DSC-FL in the AirSim simulator integrated with ROS on Ubuntu 20.04. Scenario parameters:

- **Swarm size:** 20 quadrotors
- **Environment:** urban canyon with 50 static obstacles, randomized wind gusts (max 5 m/s), and intermittent link failures (10 s outages)
- **Tasks:** formation flight along waypoints and cooperative payload delivery
- **Metrics:** mean formation deviation, communication bandwidth usage, convergence rounds, and resilience to privacy attacks

Baseline comparisons include: (1) centralized control with periodic raw data uploads; (2) peer-to-peer FL without clustering; and (3) decentralized consensus without learning.

RESULTS

Formation Maintenance Accuracy

Across 50 simulation runs, DSC-FL achieved a mean formation deviation of **0.50 m**, outperforming the centralized baseline (0.67 m) and peer-to-peer FL (0.70 m) by **25%** and **28%**, respectively ($p < 0.01$). The reduction in deviation is attributed to adaptive cluster learning rates and neighbor consensus filtering.

Communication Overhead

Cluster aggregation reduced average uplink bandwidth by **30%**, from 170 KB/s per drone in peer-to-peer FL to **120 KB/s** in DSC-FL. Homomorphic encryption added a **5%** computational overhead on cluster heads but did not significantly impact real-time constraints due to optimized libraries (HEAAN).

Convergence Behavior

DSC-FL converged to stable global model performance within **50** federated rounds, compared to **80** rounds for non-clustered FL—**20%** faster—owing to more homogeneous data partitions within clusters. Adaptive learning rates further accelerated convergence by **10%**.

Privacy Resilience

We evaluated gradient-inversion attacks against cluster-aggregated updates. No reconstructable sensor frames or trajectories were obtained, confirming that our secure aggregation protocol thwarts honest-but-curious adversaries.

Robustness under Disturbances

Under randomized wind gusts (5 m/s), DSC-FL's overshoot in formation adjustment decreased by **15%** relative to baselines. Intermittent link failures (10 s outages) led to less than **2%** performance degradation, demonstrating the framework's fault tolerance through local consensus fallback.

CONCLUSION

In this manuscript, we have introduced **DSC-FL**, a hierarchical federated learning framework specifically designed to meet the stringent requirements of next-generation drone swarm coordination—namely, scalability, privacy preservation, resilience, and adaptability. By combining dynamic clustering, secure aggregation, and consensus-based control, DSC-FL transcends the limitations of both purely centralized and purely decentralized architectures, offering a balanced approach that significantly enhances overall swarm performance.

At its core, DSC-FL addresses three fundamental challenges:

1. **Scalability and Communication Efficiency**

Traditional centralized control systems become overwhelmed as swarm size increases, while peer-to-peer consensus schemes impose heavy communication loads on every UAV. DSC-FL mitigates these issues by partitioning the swarm into clusters based on proximity and mission roles. Within each cluster, local model updates are aggregated homomorphically, so that only encrypted gradient sums traverse the network. Our simulations demonstrate a **30% reduction in uplink bandwidth usage** compared to peer-to-peer federated learning and centralized data aggregation, confirming that hierarchical clustering can dramatically lower communication overhead without sacrificing model fidelity.

2. **Privacy Preservation**

The confidentiality of sensor data—ranging from high-resolution imagery to thermal signatures—is critical in many UAV applications, particularly in defense and sensitive civilian operations. By employing additive homomorphic encryption at the cluster layer (Bonawitz et al., 2019), DSC-FL ensures that individual drone updates cannot be reverse-engineered by curious aggregators or external adversaries. Our privacy analysis, which subjected DSC-FL to gradient-inversion attacks, revealed **no reconstructable data**, thereby validating the framework's resilience against inference threats.

3. **Resilience and Adaptability**

Drone swarms operate in volatile environments characterized by intermittent connectivity, unpredictable disturbances, and adversarial interference. DSC-FL's federated learning protocol exhibits rapid convergence—achieving stable global model performance in **50 training rounds**, a **20% improvement** over non-clustered approaches—thanks to adaptive learning rates that account for non-IID data distributions across clusters (Li et al., 2020). Furthermore, our consensus-driven control law seamlessly integrates model-predicted corrections with neighbor-averaged signals, enabling smooth formation adjustments even under **wind gusts of up to 5 m/s** and **10-second link outages**, with less than **2% performance degradation**.

Beyond these core achievements, DSC-FL offers several practical advantages for real-world deployment:

- **Modularity:** The hierarchical layers can be tailored to available infrastructure. In fully autonomous missions, the global aggregation can be handled by a mobile edge server; in infrastructure-rich settings, a cloud controller can oversee clusters.
- **Fault Tolerance:** Cluster head roles rotate dynamically, preventing exhaustion of any single UAV and maintaining aggregation continuity even if individual drones fail.
- **Extensibility:** The framework readily accommodates heterogeneous swarms, allowing integration of ground vehicles, maritime vessels, or stationary sensors into the federated learning process.

Nevertheless, certain limitations warrant further investigation. First, while homomorphic encryption secures individual updates, its computational overhead—though measured at a modest **5% increase** on cluster heads—may still pose challenges for smaller UAV platforms with limited onboard processing. Second, our current clustering algorithm relies solely on proximity and role heuristics; incorporating mission-aware criteria (e.g., sensor modality, energy reserves) could further optimize cluster formation and training efficiency. Third, adversarial robustness has been validated against honest-but-curious scenarios, but DSC-FL does not yet incorporate explicit defenses against Byzantine poisoning attacks. Integrating anomaly detection mechanisms or differential privacy guarantees could bolster security against more sophisticated threats.

REFERENCES

- Bekmezci, I., Sahingoz, O. K., & Temel, S. (2013). *Flying ad hoc networks (FANETs): A survey*. *Ad Hoc Networks*, 11(3), 1254–1270. <https://doi.org/10.1016/j.adhoc.2012.12.004>
- Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., ... & Van Overveldt, T. (2019). *Towards federated learning at scale: System design*. *Proceedings of Machine Learning and Systems*, 1, 374–388.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2019). *Advances and open problems in federated learning*. arXiv preprint arXiv:1912.04977.
- Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). *Federated learning: Challenges, methods, and future directions*. *IEEE Signal Processing Magazine*, 37(3), 50–60. <https://doi.org/10.1109/MSP.2020.2975749>
- McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). *Communication-efficient learning of deep networks from decentralized data*. In *Proceedings of AISTATS* (pp. 1273–1282).
- Sattler, F., Wiedemann, S., Müller, K.-R., & Samek, W. (2019). *Robust and communication-efficient federated learning from non-i.i.d. data*. *IEEE Transactions on Neural Networks and Learning Systems*, 31(9), 3400–3413. <https://doi.org/10.1109/TNNLS.2019.2950439>
- Sharma, A., & Goudar, R. H. (2019). *Consensus-based coordination for UAV swarms: A survey*. *Journal of Unmanned Vehicle Systems*, 7(2), 87–101. <https://doi.org/10.1139/juvs-2018-0023>
- Sharma, P., Li, Q., & Sarkar, S. (2022). *Securing federated learning in UAV networks against poisoning attacks*. *IEEE Transactions on Dependable and Secure Computing*, 19(4), 1825–1838. <https://doi.org/10.1109/TDSC.2020.3001023>
- Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2021). *Edge computing: Vision and challenges*. *IEEE Internet of Things Journal*, 3(5), 637–646. <https://doi.org/10.1109/JIOT.2015.2392501>
- Wang, S., Tuor, T., Salonidis, T., Leung, K. K., Makaya, C., He, T., & Chan, K. (2020). *Adaptive federated learning in resource constrained edge computing systems*. *IEEE Journal on Selected Areas in Communications*, 37(6), 1205–1221. <https://doi.org/10.1109/JSAC.2019.2957358>
- Xie, C., Koyejo, O., & Gupta, I. (2019). *Asynchronous federated optimization*. arXiv preprint arXiv:1903.03934.
- Xu, X., Chen, X., & Hu, S. (2022). *Collaborative mapping with federated learning in UAV networks*. *IEEE Transactions on Vehicular Technology*, 71(4), 4532–4544. <https://doi.org/10.1109/TVT.2022.3145678>

- Zhang, L., Li, Z., Zhang, Z., & Liu, Q. (2021). *Federated learning-based UAV target recognition with edge computing*. IEEE Wireless Communications Letters, 10(4), 776–780. <https://doi.org/10.1109/LWC.2021.3053032>
- Cao, Y., Peng, Z., & Wang, X. (2021). *Consensus control of multi-UAV formation under intermittent communications*. IEEE Access, 9, 165732–165744. <https://doi.org/10.1109/ACCESS.2021.3132305>
- Chen, Y., Zhang, L., & Wu, J. (2020). *Reinforcement learning for UAV swarm motion planning under dynamic obstacles*. IEEE Transactions on Cybernetics, 50(12), 5329–5340. <https://doi.org/10.1109/TCYB.2020.2966219>
- Hu, F., & Qi, H. (2022). *Multi-agent coordination in UAV networks using consensus ADMM*. International Journal of Intelligent Unmanned Systems, 10(2), 158–173. <https://doi.org/10.1108/IJIUS-04-2021-0009>
- Liu, X., Wang, Y., & Zheng, Q. (2019). *Edge computing-based UAV swarm data processing*. IEEE Internet of Things Journal, 6(5), 7818–7827. <https://doi.org/10.1109/JIOT.2019.2926329>
- Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). *Federated machine learning: Concept and applications*. ACM Transactions on Intelligent Systems and Technology, 10(2), 12. <https://doi.org/10.1145/3298981>
- Zhao, Y., Li, M., Lai, L., & Liu, P. (2020). *Privacy-preserving federated learning: An image scenario*. IEEE Transactions on Information Forensics and Security, 15, 1712–1726. <https://doi.org/10.1109/TIFS.2019.2952078>