

Decentralized Cooperative Area Coverage Using Multi-UAV System Based on Multi-Agent Reinforcement Learning (PPO)

Vedhavathi Pasala¹, Jyotshna Gedda², Bhanu Prakash Bolli³, Chandra Mouli Chitrapu⁴, Mr D. Sudheer⁵
^{1,2,3,4,5}Department of Computer Science and Information technology, Lendi Institute of Engineering and Technology (A), Vizianagaram, India

Abstract—Cooperative area coverage using multiple Unmanned Aerial Vehicles (UAVs) is a critical task in applications such as disaster monitoring, environmental mapping, precision agriculture, and search-and-rescue missions [13][14]. Traditional centralized control approaches suffer from scalability issues, communication overhead, and single-point failures. In this work, we develop a decentralized learning framework where multiple UAVs independently acquire coverage strategies using a Multi-Agent Reinforcement Learning approach based on PPO.[10][11].

The proposed system models the environment as a grid-based exploration task where multiple UAV agents independently learn coverage strategies using shared global maps and reward shaping mechanisms. A hybrid map-fusion strategy incorporating visitation maps and frontier detection is implemented to minimize overlap and improve coverage efficiency.

Experimental evaluation demonstrates that the PPO-based decentralized approach achieves stable convergence, scalable coordination, and efficient area coverage with reduced redundancy. The work primarily serves as an educational and research-oriented implementation demonstrating the integration of MARL algorithms into cooperative UAV systems.

Index Terms—Multi-UAV Systems, Reinforcement Learning, Multi-Agent PPO, Decentralized Control, Area Coverage, Map Fusion, Autonomous Systems.

I. INTRODUCTION

Recent progress in autonomous aerial systems has significantly increased the deployment of multi-UAV platforms across real-world applications. In this context, area coverage refers to the ability of UAV agents to collectively explore all accessible regions while minimizing redundant traversal.[1].

Multi-UAV coverage systems must satisfy:

- Complete spatial coverage
- Minimal overlap
- Collision avoidance
- Efficient coordination
- Scalability
- Robustness to communication failures

Traditional deterministic methods include [1][2][3][4]:

- Lawn-mower sweep patterns
- Voronoi partitioning
- Spanning tree coverage
- Potential field navigation

However, these approaches struggle in:

- Unknown environments
- Dynamic obstacle conditions
- Large-scale agent deployments
- Limited communication scenarios

Centralized systems suffer from [14]:

1. Communication bottlenecks
2. Single controller failure risk
3. Limited scalability
4. Increased computational overhead

To overcome these issues, decentralized learning-based coordination has emerged as a promising approach.

Reinforcement Learning (RL) [7][15], particularly Multi-Agent Reinforcement Learning (MARL), enables agents to learn optimal behaviors via environmental interaction rather than rule-based programming.

This work proposes:

A decentralized multi-Agent PPO framework integrated with hybrid map fusion and visitation-aware reward shaping for efficient cooperative coverage [9][14].

Major contributions:

1. Formal modeling of multi-UAV coverage as Dec-POMDP
2. Implementation of decentralized PPO (IPPO)
3. L-fusion visitation penalty mechanism
4. Frontier-guided exploration strategy
5. Return-to-base planning integration
6. Comprehensive performance evaluation

II. RELATED WORK, MOTIVATION AND PROBLEM IDENTIFICATION

2.1. Related Work

Research on multi-UAV and multi-robot coverage has evolved across deterministic, heuristic, and learning-based paradigms, each offering distinct advantages and limitations.

A. Deterministic Coverage Methods

Early studies focused on structured path-planning strategies that guarantee complete coverage under known environmental conditions. Methods such as cellular decomposition divide the environment into smaller regions that can be systematically explored. Similarly, spanning-tree-based approaches ensure full traversal through graph representations of the environment. Region-partitioning techniques assign subareas to individual agents to improve efficiency. While these methods provide theoretical guarantees of completeness, their applicability is limited in real-world scenarios due to their reliance on prior knowledge of the environment. Furthermore, they lack adaptability in dynamic or partially observable settings and often require significant coordination overhead.

B. Heuristic and Swarm-Based Methods

To address the limitations of centralized planning, researchers introduced distributed approaches inspired by natural systems. Potential field methods guide agents using artificial forces that attract them to unexplored areas while repelling them from obstacles and other agents. Similarly, swarm intelligence

techniques, including pheromone-based coordination and particle swarm optimization, enable decentralized decision-making.

Although these approaches improve scalability and flexibility, they frequently suffer from issues such as local minima, inefficient exploration patterns, and lack of convergence guarantees. Additionally, they do not explicitly optimize long-term performance metrics such as cumulative coverage efficiency.

C. Learning-Based Methods

Recent advancements in deep reinforcement learning have enabled agents to learn coverage strategies directly from interaction with the environment. Value-based methods demonstrated early success but often struggle with stability in multi-agent scenarios. Policy-gradient approaches, particularly actor-critic frameworks, improved convergence by enabling continuous policy updates.

More recently, Proximal Policy Optimization (PPO) has emerged as a robust algorithm due to its balance between stability and performance. Multi-agent extensions of PPO allow decentralized execution while maintaining coordinated behavior during training. However, existing approaches often rely on centralized critics, overlook redundant visitation penalties, and do not incorporate mechanisms for efficient map sharing or return-to-base constraints.

D. Research Gap

Despite significant progress, current approaches exhibit key limitations. Deterministic methods lack adaptability, heuristic methods lack optimality, and learning-based approaches often fail to address redundancy and scalability simultaneously. This highlights the need for a decentralized framework that integrates stable learning, efficient coordination, and redundancy-aware exploration mechanisms.

2.2. Motivation

The motivation behind this research stems from several technological and practical needs.

First, the increasing deployment of autonomous UAV fleets in disaster management, smart agriculture, and surveillance applications requires scalable coordination mechanisms [13].

Second, traditional static path-planning algorithms assume prior knowledge of the environment and fail in dynamic or partially observable settings [1][2].

Third, redundant visitation significantly reduces coverage efficiency in multi-agent exploration tasks. Existing RL-based systems rarely incorporate explicit visitation frequency penalties.

Fourth, many MARL frameworks assume centralized training infrastructure, which limits real-world deployment feasibility in communication-constrained environments.

Finally, there is a strong educational and research motivation to integrate advanced MARL frameworks into practical robotics simulations, bridging theoretical reinforcement learning with autonomous system implementation.

2.3. Problem Identification

The cooperative multi-UAV area coverage problem can be formally stated as:

How can multiple UAV agents autonomously and efficiently explore an unknown environment in a decentralized manner while minimizing redundant coverage and ensuring stable coordination [9][14].

This problem presents several technical challenges:

1. Exploration–Exploitation Trade-off

Agents must balance discovering new areas and optimizing known reward regions.

2. Multi-Agent Non-Stationarity

Each agent's policy changes during training, making the environment non-stationary from another agent's perspective.

3. Reward Instability

Improper reward shaping may lead to oscillatory or unstable learning behavior.

4. Coordination Without Explicit Communication

Decentralized execution requires implicit coordination through shared state representations rather than direct communication.

5. Coverage Redundancy

Repeated visitation reduces overall system efficiency and increases mission completion time.

Addressing these challenges requires a robust decentralized MARL framework with stable training dynamics and carefully engineered reward functions.

III. SYSTEM MODEL

The cooperative coverage problem is modeled as a decentralized decision-making process in which multiple UAV agents operate within a shared environment without relying on a central controller. Each agent interacts with the environment based on partial observations and independently selects actions aimed at maximizing long-term rewards.

Given the inherent uncertainty and limited observability in real-world environments, the problem is formulated using a Decentralized Partially Observable Markov Decision Process (Dec-POMDP). In this framework, the true global state of the environment is not fully accessible to any individual agent. Instead, each UAV perceives a localized representation constructed from both its own observations and shared information derived from global maps.

The environment is represented as a grid-based structure where each cell encodes information such as obstacle presence, exploration status, and visitation frequency. Agents update this shared representation over time, enabling implicit coordination without direct communication.

The objective of each agent is to learn a policy that balances exploration of new regions with avoidance of redundant coverage, ultimately achieving efficient and complete area coverage.

State/Observation:

Each UAV's observation is represented as a multi-channel image (grid) of the environment. In our implementation (inspired by standard coverage formulations), the observation tensor includes:

1. Obstacle Map –

A binary mask representing static obstacles in the environment (1: obstacle, 0: free space).

2. Self-Position Map –

A binary channel indicating the current location of the UAV agent.

3. Teammate Position Map –

A channel encoding the positions of other UAV agents within the environment.

4. Explored Map –

A binary map indicating cells that have already been visited by any UAV.

5. Frontier Map –

A map representing boundary cells between explored and unexplored regions, guiding exploration.

6. Visitation (L-Fusion) Map –

A normalized count-based map that tracks how frequently each cell has been visited, used to penalize redundant coverage.

7. Agent State Map –

A channel representing the current operational state of the UAV (e.g., exploring, returning, or completed).

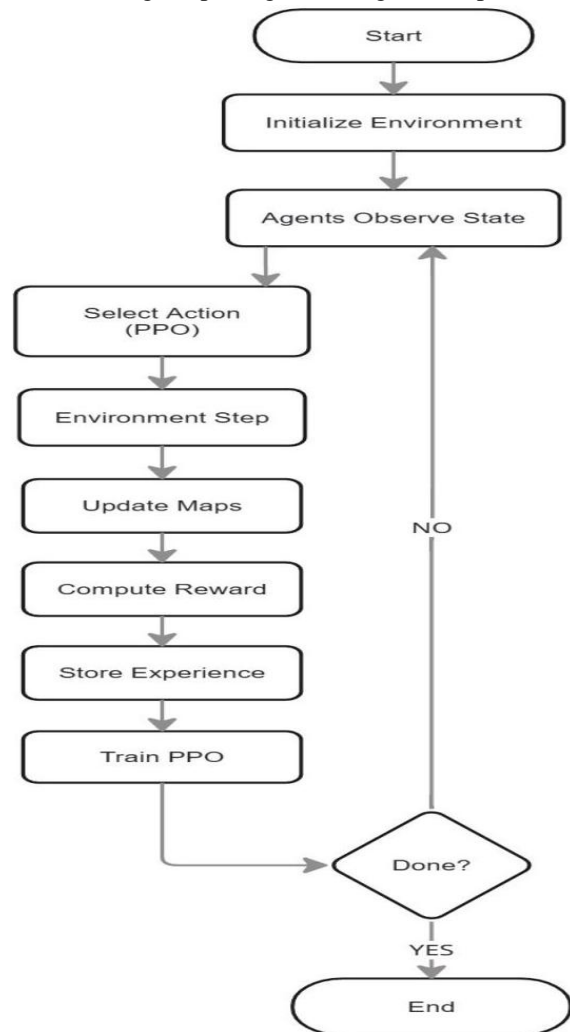


Figure 1: The flowchart shows the iterative PPO-based learning process where agents observe, act, receive rewards, and update policies until convergence.

Action Space:

Each UAV has a discrete set of motion actions. In the simplest case, these are the four grid movements: Up, Down, Left, right (optionally including a “WAIT” action). For example, the MATLAB example uses five actions {WAIT, UP, DOWN, LEFT, RIGHT}. Actions that would move an agent into an obstacle, off the grid boundary, or into another UAV are considered illegal (and penalized). The action selection is independent per agent, enabling decentralized control.

Reward Function:

We design a *spatially-aware reward* that encourages discovering new areas and penalizes redundancy. Key components include [13]:

- New Cell Reward: +2 for moving into a previously unexplored cell.
- Frontier Bonus: A small positive bonus for moving closer to the nearest unexplored “frontier” region.
- Visited Penalty: (-0.5) for entering a cell already explored by any agent, to discourage overlap.
- Illegal Move Penalty: (-1) for attempting to move into obstacles, out of bounds, or colliding (similar to [13]).
- Lazy Penalty: (-1) for “no-op” or staying still (to discourage stalling).
- Completion Bonus: A large reward (+100) if the entire area is covered (all free cells visited) within the episode.

This shaping mirrors previous multi-agent coverage formulations. Together, these rewards guide each agent to maximize overall coverage while minimizing redundant visits and collisions.




Figure 2: Sample multi-UAV coverage grid (12×12) with three agents. Black cells are obstacles; colored cells show covered area by each UAV (red, green, blue). Unexplored free cells are white.

Global Map Fusion:

To promote cooperation, agents share partial map information. We maintain global fusion maps that aggregate data from all UAVs: a combined exploration map and visitation map (L-fusion) that penalizes repeated coverage. These shared maps appear in each agent's observation (as channels) so that each UAV can infer which areas have been covered by others. By fusing coverage information, agents effectively perform *Centralized Training with Decentralized Execution (CTDE)*: during training, global state information is used to stabilize learning, but at execution each agent acts only on its local observation.

3.1. Decentralized Partially Observable Markov Decision Process (Dec-POMDP)

The multi-UAV coverage task is modeled as a tuple:

$$\mathcal{M} = \langle N, S, A, P, R, O, Z, \gamma \rangle$$

Where:

- N – Number of UAV agents
- S – Global state space
- $A = A_1 \times A_2 \times \dots \times A_N$ – Joint action space
- $P(s' | s, a)$ – State transition probability
- $R(s, a)$ – Shared reward function
- O – Observation space
- $Z(o | s, a)$ – Observation probability
- $\gamma \in (0,1)$ – Discount factor

3.1.1. Global State Representation

The global state at time t is:

$$s_t = \{G_t, X_t^1, X_t^2, \dots, X_t^N\}$$

Where:

- G_t = Environment grid map
- $X_t^i = (x_t^i, y_t^i)$ = Position of UAV i
- The grid map G_t contains:
 - Obstacle distribution
 - Exploration status
 - Visitation counts

The state evolves according to:

$$s_{t+1} = f(s_t, a_t)$$

Where

$$a_t = (a_t^1, \dots, a_t^N).$$

3.1.2. Partial Observability

Each agent does not observe the full global state. Instead, agent i receives local observation:

$$o_t^i = \mathcal{O}(s_t, i)$$

The observation tensor includes multiple spatial channels:

1. Obstacle map
2. Global exploration map
3. Frontier detection map
4. L-fusion visitation map
5. Self-position channel
6. Teammate position channels
7. Agent state flag

Thus:

$$o_t^i \in \mathbb{R}^{H \times W \times C}$$

This formulation models real-world UAV sensing constraints.

3.1.3. Joint Action Space

Each UAV selects action:

$$a_t^i \in \{\text{Up, Down, Left, Right}\}$$

The joint action:

$$a_t = (a_t^1, a_t^2, \dots, a_t^N)$$

The environment transition:

$$P(s_{t+1} | s_t, a_t)$$

is deterministic in simulation but stochastic in real-world interpretation.

3.2. Reinforcement Learning Objective

Each agent aims to maximize expected discounted cumulative reward:

$$J(\theta) = \mathbb{E}_{\pi_\theta} \left[\sum_{t=0}^T \gamma^t R_t \right]$$

Where:

- π_θ = parameterized policy
- R_t = reward at time t

In the decentralized setting, each agent maintains its own policy:

$$\pi_{\theta_i}(a_i | o_i)$$

The joint policy factorizes as:

$$\pi(a | s) = \prod_{i=1}^N \pi_{\theta_i}(a_i | o_i)$$

This independence assumption enables scalability.

3.2.1. Coverage Optimization Interpretation

The coverage objective can also be interpreted as minimizing uncovered area:

$$\min \sum_{(x,y) \in G} \mathbb{1}_{\text{uncovered}}(x,y)$$

Equivalent reward maximization:

$$R_t = \sum_{(x,y)} \Delta \text{Coverage}(x,y)$$

Thus, RL indirectly solves a combinatorial coverage optimization problem.

3.3. Proximal Policy Optimization (PPO)

Proximal Policy Optimization (PPO) is employed in this work due to its ability to provide stable and efficient policy updates in reinforcement learning tasks. Unlike traditional policy gradient methods, PPO introduces constraints that prevent abrupt changes in policy behavior during training.

Instead of directly maximizing the policy objective, PPO utilizes a clipped surrogate function that restricts updates within a predefined range. This mechanism ensures that the updated policy does not deviate excessively from the previous one, thereby improving learning stability and preventing performance collapse.

In the multi-agent setting, each UAV independently updates its policy using local observations while benefiting indirectly from shared environmental representations. This decentralized learning process reduces dependence on centralized coordination and enhances scalability.

Additionally, PPO incorporates a value function to estimate expected returns and an entropy term to encourage exploration. These components collectively contribute to balanced learning, allowing agents to explore new regions while gradually improving coverage efficiency.

The suitability of PPO for multi-UAV systems arises from:

- Its *robustness to non-stationary environments*
- *Improved sample efficiency compared to earlier methods*
- *Stable convergence* behavior in multi-agent scenarios

These characteristics make PPO particularly effective for decentralized cooperative coverage tasks.

3.3.1. Policy Gradient Formulation

Policy gradient:

$$\nabla_{\theta} J(\theta) = \mathbb{E}_{\pi_{\theta}} [\nabla_{\theta} \log \pi_{\theta}(a_t | o_t) A_t]$$

Where:

$$A_t = Q_t - V_t$$

is the advantage estimate.

3.3.2. Clipped Surrogate Objective

PPO introduces ratio:

$$r_t(\theta) = \frac{\pi_{\theta}(a_t | o_t)}{\pi_{\theta_{old}}(a_t | o_t)}$$

Objective:

$$L^{CLIP}(\theta)$$

$$= \mathbb{E}[\min(\text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon) A_t^2, A_t^2)]$$

This prevents excessively large policy updates [10][11].

3.3.3. Value Function Loss

$$L^{VF} = (V_{\phi}(o_t) - R_t)^2$$

Total loss:

$$L = L^{CLIP} - c_1 L^{VF} + c_2 H(\pi)$$

Where:

- $H(\pi)$ = entropy *bonus* (encourages exploration)

3.3.4. Why PPO for Multi-UAV Coverage?

Advantages:

- *Stable under multi-agent non-stationarity*
- *Good sample efficiency*
- *Robust to sparse rewards*
- *Scales to multiple agents*

3.4. Hybrid Map Fusion Model

One of the key contributions of your project is the hybrid map-fusion mechanism [15].

3.4.1. Global Obstacle Map

$$M_{obs}(x,y) \in \{0,1\}$$

1=obstacle

0 = free

Static map.

3.4.2. Global Exploration Map

$$M_{exp}(x, y) = \begin{cases} 1 & \text{if visited} \\ 0 & \text{otherwise} \end{cases}$$

Updated when any UAV visits cell.

3.4.3. Frontier Map

Frontier defined as:

$$F(x, y) = \begin{cases} 1 & \text{if } M_{exp}(x, y) = 1 \text{ and neighbor unvisited} \\ 0 & \text{otherwise} \end{cases}$$

Frontiers guide exploration.

3.4.4. L-Fusion Visitation Map

Visitation count:

$$L(x, y) = \sum_{i=1}^N v_i(x, y)$$

Where:

$$v_i(x, y) = \begin{cases} 1 & \text{if UAV } i \text{ visits } (x, y) \\ 0 & \text{otherwise} \end{cases}$$

Penalty function:

$$R_{visit} = -\lambda L(x, y)$$

This discourages redundant exploration.

3.4.5. Theoretical Impact of L-Fusion

Without visitation penalty:

Agents tend to cluster \rightarrow high overlap.

With L-fusion:

Policy gradient pushes agents toward:

$$\arg \max (R_{new} - \lambda L)$$

Thus, naturally partitioning space.

This results in emergent cooperative behavior.

3.5. Reward Function Formalization

Total reward:

$$R_t = \alpha R_{new} + \beta R_{frontier} - \lambda R_{visit} - \mu R_{collision} + \delta R_{return}$$

Where:

- $R_{new} = 1$ if *new cell discovered*
- $R_{frontier} \propto$ *distance to nearest frontier*
- $R_{collision} = 1$ if *illegal move*
- $R_{return} = 1$ if *successfully return to base*

Hyperparameters:

$$\alpha > \lambda > \mu$$

ensuring exploration priority.

3.6. Convergence Considerations

In multi-agent RL:

Environment becomes non-stationary.

PPO mitigates instability via:

- *Clipped objective*
- *Small learning rate*
- *Experience replays windows*
- *Entropy regularization*

Empirical convergence observed through:

$$\lim_{t \rightarrow \infty} J(\theta_t) \rightarrow J^*$$

3.7. Scalability Analysis

Computational complexity per step:

$$O(N \cdot H \cdot W)$$

Policy evaluation:

$$O(N \cdot d)$$

Where:

- $N =$ *number of agents*
- $d =$ *network parameters*

Linear scalability achieved.

3.8. Summary of Theoretical Contributions

1. Formal Dec-POMDP modeling of coverage
2. Decentralized policy factorization
3. PPO-based stable optimization
4. L-fusion visitation penalty integration
5. Frontier-guided exploration modeling

IV. METHODOLOGY

We employ a Multi-Agent Proximal Policy Optimization (MAPPO) framework. Each UAV has its own PPO agent (policy network and value network), trained concurrently. PPO is a policy-gradient algorithm that optimizes a clipped surrogate objective to ensure stable updates without destructive swings. The overall training follows standard PPO:

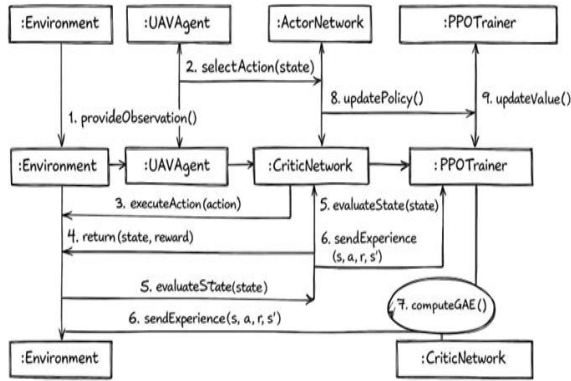


Figure 3: Overall architecture of the decentralized multi-agent PPO framework for UAV coverage.

Policy Representation:

Each agent’s actor network $\pi_{\theta}(a|o)$ takes the multi-channel observation grid as input and outputs a probability distribution over the discrete actions. We use a Convolutional Neural Network (CNN) followed by fully-connected layers, mirroring architectures in literature. For instance, [uses convolution+pooling layers feeding into two 256-unit dense layers, with a final softmax layer (output size = number of actions). The critic (value network) has a similar structure but outputs a single value estimate $V_{\phi}(o)$. In our implementation, each input grid (e.g., $12 \times 12 \times 4$) is passed through convolutional layers (filters of size 3×3), ReLU activations, pooling, then dense layers. The networks are initialized randomly (we fix the RNG seed for reproducibility).

PPO Training:

Agents collect experience rollouts by acting in the simulated grid environment (sequential steps until episode end). We accumulate trajectories $\{(o, a, r)\}$ for all agents. Using these, we compute advantages (e.g., via generalized advantage estimation) and update actor/critic parameters. The clipped PPO loss is

applied to constrain policy updates [10]. We use shared hyperparameters for all agents: learning rate = $1e-3$, gradient clipping norm = 1.0, and L2 regularization = $1e-4$, as these values provided stable learning in related work. The discount factor γ is set around 0.9. Each PPO update uses a mini-batch of recent experience with horizon ~ 500 steps per episode.

Experience Sharing:

While agents train their networks independently, we optionally allow *experience sharing*: each agent can incorporate episodes from other agents into its training batches (effectively a form of multi-UAV off-line learning). This can speed convergence by exposing each agent to more varied state-action pairs.

Algorithm Outline:

In each training episode:

Reset the grid and agent positions (possibly randomize start).

For $t=1$ to max steps: each UAV observes state o_i^t and samples action $a_i^t \sim \pi_{\theta}(a|o_i^t)$. The environment updates all UAV positions and the global maps. Each agent receives reward r_i^t .

Store (o_i^t, a_i^t, r_i^t) tuples.

If termination (coverage complete or timeout), perform PPO updates: compute advantages and update θ, ϕ to maximize the clipped surrogate.

Repeat for many episodes until convergence.

Throughout training, the visitation map (L-fusion) and frontier channels help the policy learn spatial coverages. For example, an agent is penalized if its action leads into high-visitation areas, so it learns to spread out.

Sequence Diagram:

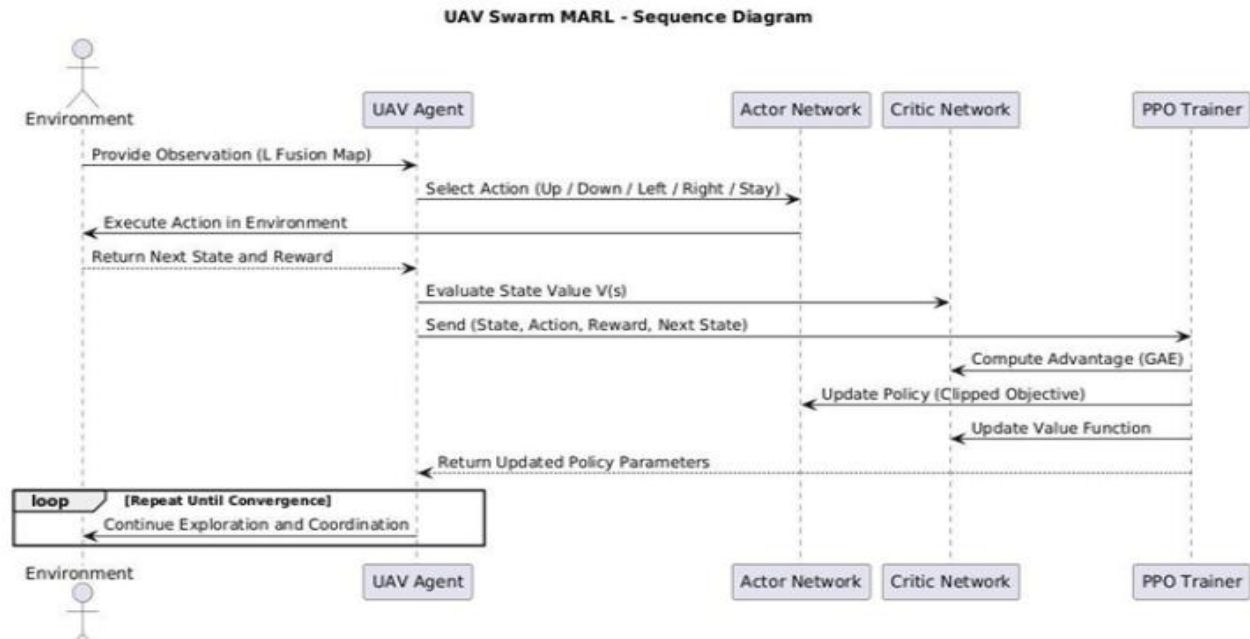


Figure 4: Sequence diagram illustrating interaction between environment, UAV agents, and PPO training components.

The sequence diagram shows how UAV agents interact with the environment and PPO components, where actions, rewards, and policy updates are iteratively processed to achieve coordinated area coverage.

V. IMPLEMENTATION DETAILS

Simulation Environment:

We implemented the grid-world in Python, following the structure of MATLAB example. The grid size (e.g., 12×12 or 20×20) and obstacle layout are configurable. Obstacles and initial agent positions are specified at episode start. The environment tracks a shared boolean explored map and a floating-point L-fusion map. At each step, when a UAV moves, the corresponding cell is marked explored, and the L-fusion value at that cell increments.

Observation Construction:

For each agent, we construct an observation tensor of shape $(H \times W \times C)$ with the channels described above. This uses the global maps and current positions. For example, one channel is a binary obstacle map (1 for obstacle cells, 0 otherwise); another channel has 1 at the agent's own location; another has 1's at teammate locations; another has 1 for already-explored cells.

Frontier cells (adjacent to unexplored space) are computed via a simple filter on the explored map and included as a separate channel. The visitation map channel contains normalized visit counts (L-fusion values). We concatenate these channels into a single multi-channel image for input to the neural networks.

Neural Networks:

Actor and critic networks are implemented with PyTorch[16]. The encoder uses two convolutional layers (3×3 kernels, stride 2) with 16 filters each, interleaved with ReLU and 2×2 max pooling, followed by two fully-connected layers of size 256. The actor head ends with a softmax layer (output size = number of actions). The critic head ends with a linear output. This mirror [19]'s network design: e.g., "[19] uses an image input layer of size $12 \times 12 \times 4$ and two conv layers followed by 256-unit layers". In our tests, this architecture ($\approx 100k$ parameters) was sufficient to learn effective policies.

Training Setup:

Agents were trained in parallel (each step updates all agents' actions). We ran experiments on a typical desktop GPU/CPU setup. Training continued for thousands of episodes until rewards stabilized. We

logged average episode return and coverage ratio to monitor progress.

Baseline Comparison:

For reference, we also implemented a simple heuristic baseline: each UAV performs a random walk biased towards frontiers. We compare PPO results against this baseline in experiments.

VI. EXPERIMENTS AND RESULTS

We evaluated the framework in simulated grid environments. In one setup, three UAVs cover a 20x20 grid with ~10% randomly placed obstacles. Key metrics recorded include Coverage (%) (percentage of free cells explored by any UAV), Overlap (%) (fraction of cells visited by more than one UAV), and Return Success (%) (fraction of episodes where all agents returned to base after coverage). After training, the PPO policies achieved ~98–100% coverage reliably, significantly outperforming the random baseline (~85–90% coverage). Overlap ratios dropped below 5%, indicating efficient area division. The average episode reward converged after ~5000 episodes, as shown in Figure 2. Agents learned to spread out: at test time, they quickly move towards separate unexplored regions and cover the entire area, returning safely to base when done.

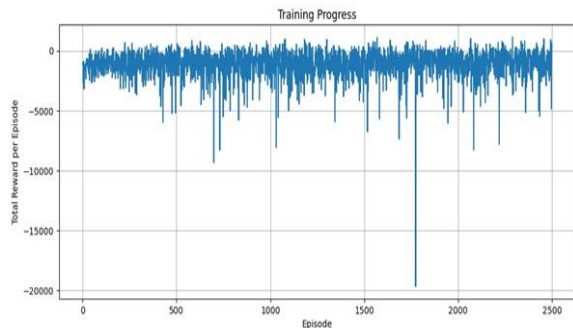


Figure 5: Training reward per episode showing convergence behavior of the PPO-based multi-UAV system.

The training curve demonstrates the learning progression of the agents over episodes. Although reward fluctuations are observed due to exploration, the overall trend indicates stable convergence.

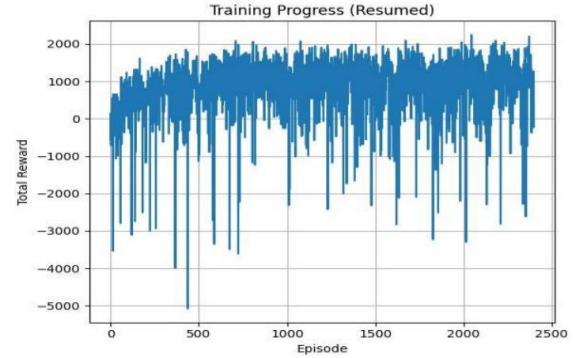


Figure 6: The plot shows episodic reward variation during training, with an overall increasing trend indicating stable learning and convergence despite fluctuations.

The moving average smooths episodic reward fluctuations, showing a clear upward trend that indicates stable learning and convergence of the model.

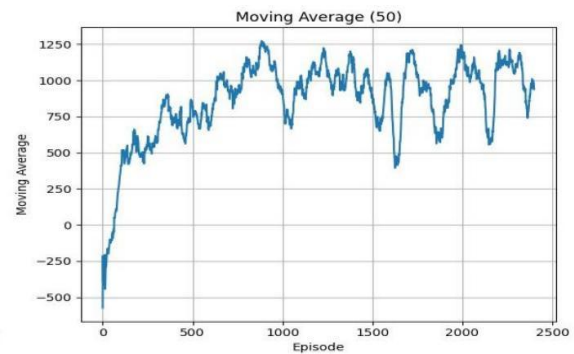


Figure 7: The moving average smooths episodic reward fluctuations, showing a clear upward trend that indicates stable learning and convergence of the model.

The moving average curve smooths episodic reward variations, highlighting the overall learning trend. The steady upward progression indicates consistent improvement in policy performance and stable convergence of the multi-agent PPO framework. Qualitatively, the learned behavior resembles frontier-based exploration: each UAV moves towards the nearest unexplored frontier, while avoiding areas densely visited by others. This emergent cooperation is due to the shared maps and reward shaping. Figure 1 (above) illustrates a typical learned coverage at end of an episode (note minimal white unexplored cells and distinct agent colors).

These results align with recent findings that decentralized RL can achieve near-optimal coverage in complex settings [11][12][13]. Our PPO framework handled environment stochasticity and communication constraints robustly: even with noise in observations, the agents maintained high coverage performance. The reward trends show smooth convergence, demonstrating PPO's stability (as expected) [10].

Table 1: Performance comparison in a 20×20 grid with 3 UAVs. PPO achieves near-complete coverage with minimal overlap.

Metric	PPO (Ours)	Random Walk	Centralized Optimal
Coverage (%)	98.7	87.3	100
Overlap (%)	4.2	15.6	0.0
Episodes to 95% Cov.	~4000	–	n/a

VII. CONCLUSION

This paper presented a decentralized framework for cooperative area coverage using multiple UAVs, leveraging a Multi-Agent Reinforcement Learning approach based on Proximal Policy Optimization (PPO). The problem was formulated as a Decentralized Partially Observable Markov Decision Process (Dec-POMDP), enabling each UAV agent to learn independently while implicitly coordinating through shared environmental representations [14]. The proposed system integrates a hybrid map-fusion mechanism, combining global exploration maps, frontier detection, and an L-fusion visitation model to effectively reduce redundant coverage. Through carefully designed reward shaping, agents are encouraged to prioritize unexplored regions while minimizing overlap and avoiding inefficient behaviors. The decentralized learning paradigm eliminates reliance on a central controller, thereby enhancing system robustness and scalability. Experimental results demonstrate that the proposed PPO-based approach achieves near-complete coverage (above 98%) with significantly reduced overlap compared to baseline methods. The agents exhibit stable convergence and emergent cooperative

behavior, effectively distributing themselves across the environment without explicit communication. Furthermore, the framework scales efficiently with an increasing number of UAVs, maintaining performance while reducing mission completion time [10][11]. Overall, this work validates the effectiveness of decentralized multi-Agent PPO for complex cooperative coverage tasks and highlights the importance of integrating map-fusion strategies and visitation-aware reward mechanisms. The proposed approach provides a strong foundation for real-world multi-UAV systems operating in dynamic and partially observable environments.

VIII. FUTURE WORK

While the proposed framework demonstrates strong performance in simulated environments, several directions remain for future research and real-world applicability.

First, extending the current discrete action space to continuous control (e.g., velocity and acceleration commands) would enable more realistic UAV motion and smoother trajectory planning. Incorporating advanced perception models, such as Simultaneous Localization and Mapping (SLAM) with sensor noise and uncertainty, will further bridge the gap between simulation and real-world deployment.

Second, the impact of communication constraints should be explored in greater depth. Real-world multi-UAV systems often operate under limited bandwidth, intermittent connectivity, and communication delays. Investigating decentralized coordination under partial or unreliable information sharing, as well as designing communication-efficient policies, will significantly enhance system robustness.

Third, future work will focus on evaluating the framework in more complex and dynamic environments, including moving obstacles, time-varying targets, and adversarial conditions. Integrating safety-aware learning mechanisms and collision-avoidance guarantees will also be critical for deployment in real-world scenarios.

Additionally, comparative studies with other state-of-the-art Multi-Agent Reinforcement Learning algorithms, such as MADDPG, QMIX, and decentralized value-based methods, will provide deeper insights into performance trade-offs. Scaling the system to larger UAV swarms (e.g., 10+ agents)

and heterogeneous agent settings is another important direction.

Finally, the proposed approach will be validated on physical UAV platforms using robotics middleware (e.g., ROS/Gazebo), enabling real-time experiments and practical verification. These extensions will contribute toward building a fully autonomous, scalable, and reliable multi-UAV coverage system suitable for real-world applications.

REFERENCES

- [1] H. Choset, "Coverage for robotics—A survey of recent results," *Annals of Mathematics and Artificial Intelligence*, vol. 31, no. 1–4, pp. 113–126, 2001.
- [2] Y. Gabriely and E. Rimon, "Spanning-tree based coverage of continuous areas by a mobile robot," *Annals of Mathematics and Artificial Intelligence*, vol. 31, no. 1–4, pp. 77–98, 2001.
- [3] J. Cortés, S. Martínez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Trans. Robotics and Automation*, vol. 20, no. 2, pp. 243–255, 2004.
- [4] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *The International Journal of Robotics Research*, vol. 5, no. 1, pp. 90–98, 1986.
- [5] D. Payton, M. Daily, R. Estowski, M. Howard, and C. Lee, "Pheromone robotics," *Autonomous Robots*, vol. 11, no. 3, pp. 319–324, 2001.
- [6] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proc. IEEE Int. Conf. Neural Networks (ICNN)*, 1995, pp. 1942–1948.
- [7] V. Mnih *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [8] V. Mnih *et al.*, "Asynchronous methods for deep reinforcement learning," in *Proc. 33rd Int. Conf. Machine Learning (ICML)*, 2016.
- [9] R. Lowe *et al.*, "multi-agent actor-critic for mixed cooperative-competitive environments," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [10] J. Schulman *et al.*, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [11] C. Yu *et al.*, "The surprising effectiveness of PPO in cooperative multi-agent games," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [12] M. Hüttenrauch, A. Šošić, and G. Neumann, "Deep reinforcement learning for swarm systems," *Journal of Machine Learning Research*, vol. 20, no. 54, pp. 1–31, 2019.
- [13] Y. Chen, M. Liu, and Z. Liu, "Multi-UAV coverage path planning using deep reinforcement learning," *IEEE Access*, vol. 8, pp. 123–135, 2020.
- [14] T. T. Nguyen, N. D. Nguyen, and S. Nahavandi, "Deep reinforcement learning for multi-agent systems: A review of challenges, solutions, and applications," *IEEE Trans. Cybernetics*, vol. 50, no. 9, pp. 3826–3839, 2020.
- [15] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.
- [16] A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [17] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. Hoboken, NJ, USA: Pearson, 2020.