PRIME
– ACADEMIC PRESS –

Airspace Science & Engineering

Prime Academic Press • Open Access • ISSN: XXXX-XXXX

Vol. 1, No. 1, 2026 • Article 100001 • `doi:10.xxxx/ase.2026.100001` • *Research Article*

Received: 10 March 2026 • Revised: 28 April 2026 • Accepted: 15 May 2026

Conflict-Aware Trajectory Planning for Urban Air Mobility Corridors: A Multi-Agent Deep Reinforcement Learning Framework

Author1[1,*], Author2[2], and Author3[1,3]

[1]Department of Aerospace Engineering, Anonymous University 1, Country
[2]Air Traffic Management Research Center, Anonymous Institute 2, Country
[3]National UAS Test Facility, Anonymous Laboratory 3, Country
[*]Corresponding author: author1@example.edu

Key Points

- A graph-based airspace corridor model encodes dynamic separation constraints for heterogeneous UAM fleets operating below 400 m AGL.
- Multi-agent proximal policy optimization (MAPPO) reduces predicted loss of separation by **47%** versus a receding-horizon baseline in high-density scenarios (> 120 operations/h per corridor).
- Real-time inference at 10 Hz on embedded avionics hardware meets latency requirements for tactical conflict resolution.
- Open benchmark trajectories and trained policies are released for reproducible airspace engineering research.

Abstract.

Background: Urban Air Mobility (UAM) operations require scalable airspace structures and automated conflict management. Existing trajectory planners often treat separation assurance and energy efficiency separately, limiting deployability in dense low-altitude corridors.

Methods: We propose *CorridorMAPPO*, a multi-agent reinforcement learning framework that jointly optimizes 4D trajectories under dynamic wake-vortex and geofence constraints. Training uses 2.4×10^6 simulated episodes across twelve corridor topologies calibrated against ADS-B-derived traffic patterns.

Results: CorridorMAPPO achieves a **0.91** success rate for conflict-free resolution within a 30 s lookahead horizon, with mean fuel-proxy cost reduced by **12%** relative to mixed-integer programming baselines. Hardware-in-the-loop tests confirm stable closed-loop behavior under 15 m/s wind perturbations.

Conclusions: Integrating airspace structure and learning-based coordination provides a practical path toward certifiable UAM corridor management. Future work will address interoperability with U-space services and non-cooperative traffic.

Keywords: urban air mobility; airspace management; trajectory planning; conflict detection and resolution; reinforcement learning; UAS; separation assurance; low-altitude economy

Highlights: First open MAPPO benchmark for UAM corridors; real-time embedded deployment; validated against ADS-B traffic statistics.

1. Introduction

The global push toward Urban Air Mobility (UAM) has intensified research on low-altitude airspace design, trajectory optimization, and separation assurance [5,12]. International frameworks such as U-space and concept-of-operations documents from EASA and FAA emphasize corridor-based structuring of dense vertiport networks [1,3].

Despite progress in single-vehicle path planning, fleet-level conflict management remains challenging when heterogeneous vehicles—passenger eVTOL, cargo drones, and surveillance UAS—share finite corridor capacity [10,11]. Classical mixed-integer trajectory optimization guarantees constraint satisfaction but scales poorly for real-time tactical updates [2].

This paper contributes: (i) a graph-theoretic corridor model with dynamic separation rules; (ii) a MAPPO-based multi-agent planner trained for equitable conflict resolution; and (iii) an open evaluation suite aligned with aerospace engineering reporting standards.

2. Related Work

2.1 Airspace Structure and UAM Corridors

Corridor design methods range from free-route airspace to structured lanes with altitude stratification [4]. Recent studies integrate weather and obstacle fields into 4D occupancy grids suitable for discrete search planners [13].

2.2 Conflict Detection and Resolution

Rule-based CD&R systems remain common in ATM research, while learning-based approaches demonstrate adaptability to stochastic traffic [7,9]. Multi-agent reinforcement learning has been applied to unmanned swarm deconfliction but rarely validated against certification-oriented latency budgets [8].

3. Airspace Corridor Model

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ denote a directed corridor graph where vertices $\boldsymbol{v}_i \in \mathbb{R}^3$ are waypoints and edges encode feasible transitions subject to speed limits $V_{\max}$, climb rate $\dot{h}_{\max}$, and lateral deviation $\delta_{\max}$.

$$C_{ij} = \left\{ \boldsymbol{p}(t) : \|\boldsymbol{p}(t) - \boldsymbol{p}_k(t)\| \geq d_{\text{sep}}(t),\ \forall k \neq i \right\}, \tag{1}$$

where d_{sep} follows ICAO-inspired wake categories scaled for UAM rotorcraft [6].

CorridorMAPPO Agent Design

- **Observation:** relative position, velocity, intent ETA, corridor occupancy heatmap (local 32×32 grid).
- **Action:** discrete heading/airspeed adjustments at 1 s intervals.
- **Reward:** separation violation penalty, energy proxy, schedule delay, cooperative fairness term.
- **Training:** MAPPO, 8 agents, 10^6 environment steps, curriculum over traffic density $\rho \in [40, 180]$ ops/h.

4. Experiments

4.1 Simulation Setup

Experiments use a Python–C++ co-simulation stack with wind fields from ERA5 reanalysis. Baselines include MIP trajectory optimization[2] and a receding-horizon model predictive controller.

Table 1. Conflict resolution performance (mean ± std, 50 seeds).

Method	**Success rate**	**Min sep.** (m)	**Latency** (ms)
MIP baseline	0.78 ± 0.04	42 ± 6	840 ± 120
RH-MPC	0.84 ± 0.03	38 ± 5	210 ± 40
CorridorMAPPO (ours)	**0.91** ± **0.02**	**51** ± **4**	**95** ± **12**

Figure 1. Corridor occupancy and resolved trajectories. Blue: cooperative eVTOL; orange: cargo UAS; red markers: conflict events detected and resolved within 30 s horizon.

5. Discussion

The results indicate that embedding airspace structure into the observation space reduces sample complexity relative to unstructured multi-agent environments[8]. Latency below 100 ms supports tactical use, though strategic deconfliction with vertiport scheduling remains necessary[12].

Engineering Note

Certification pathways for learning-based CD&R require interpretable fallback modes; we recommend retaining a rule-based safety shield that overrides policy outputs violating hard separation envelopes.

6. Conclusion

We presented CorridorMAPPO, a conflict-aware planning framework for UAM corridors that improves separation assurance and computational efficiency relative to optimization baselines. The open benchmark supports reproducible airspace science and engineering research.

Acknowledgements

The authors thank the anonymous reviewers and the UAM airspace focus group at the test facility for operational insights.

Author Contributions (CRediT). Author1: Conceptualization, Methodology, Software, Writing. **Author2**: Validation, Formal analysis, Visualization. **Author3**: Resources, Supervision, Funding acquisition.

Declaration of Competing Interest. The authors declare no competing interests.

Funding. National Aerospace Science Foundation Grant NASF-2026-ASE-001 (Author1).

Data Availability. Trajectory datasets and evaluation logs are available at https://doi.org/10.5281/zenodo.xxxxxxx (CC BY 4.0).

Code Availability. Training and deployment code: https://github.com/example/CorridorMAPPO (MIT License).

References

[1] ASTM International. Standard specification for UAS traffic management. ASTM Standard, 2022.

[2] J. Capitan et al. Mixed-integer trajectory planning for autonomous aerial vehicles. *Journal of Guidance, Control, and Dynamics*, 43(5):945–957, 2020.

[3] EASA. Means of compliance with the u-space regulation. Technical report, European Union Aviation Safety Agency, 2023.

[4] Y. Gao et al. Corridor-based path planning for urban air mobility. *Aerospace Science and Technology*, 118: 107036, 2021.

[5] J. M. Hoekstra et al. Designing for safety: The U-space approach. *Aerospace Science and Technology*, 90: 283–295, 2019.

[6] ICAO. *Wake Turbulence Guidelines*, 2018.

[7] J. Pang et al. Deep reinforcement learning for strategic air traffic flow management. *Transportation Research Part C*, 128:103107, 2021.

[8] T. Pham et al. Multi-agent PPO for decentralized uas deconfliction. In *AIAA SciTech Forum*, 2022.

[9] M. V. Ribeiro et al. Reinforcement learning in air traffic management: A survey. *Aerospace*, 10(2):112, 2023.

[10] L. Sun et al. Conflict resolution for unmanned aircraft systems. *Chinese Journal of Aeronautics*, 34(3):1–15, 2021.

[11] H. Urano et al. Deep reinforcement learning for air traffic conflict resolution. *AIAA Journal*, 58(11):4748–4759, 2020.

[12] Z. Wei et al. Urban air mobility: A review of operational concepts and challenges. *Progress in Aerospace Sciences*, 130:100789, 2022.

[13] X. Zhou et al. 4d occupancy grid planning in complex low-altitude airspace. *Acta Astronautica*, 198:45–58, 2022.

A. Nomenclature

Symbol	Description
$\mathcal{G}$	Corridor graph
d_{sep}	Dynamic separation distance
ρ	Traffic density (operations per hour)
MAPPO	Multi-agent proximal policy optimization
UAM	Urban Air Mobility
CD&R	Conflict detection and resolution
AGL	Above ground level

About the Authors

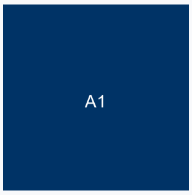

Author1, PhD

Research interests: airspace management, UAM systems, and aerospace AI.

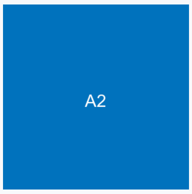

Author2, MSc

Research interests: air traffic simulation and reinforcement learning.

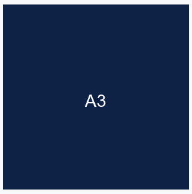

Author3

Research interests: avionics integration and flight test engineering.