

A Graph-Based Multi-Agent UAV–UGV Framework for Disaster Response and Victim Localization

Rajanikanth Javvadi¹, Shiva Shankar Reddy²

¹Faculty of Engineering and Built Science, Lincoln University College, Petaling Jaya, Selangor-47301, Malaysia; ² Department of Computer Science and Engineering, Lincoln University College, Petaling Jaya, Selangor-47301, Malaysia

Email ID: pdf.rajanikanthjavvadi@lincoln.edu.my; pdfsv.shiva@lincoln.edu.my;

Abstract: The system utilises multi-sensor perception fusion, cross-agent coordination, weather-adaptive 3D path planning, victim localization with 4-class triage prediction, and uncertainty-calibrated rescue dispatch. UAVs are used for wide-area aerial perception and thermal monitoring, and UGVs are used for payload transport and manipulation, and persistent ground operation. GUARD-X was tested on the MBZIRC 2020 benchmark and six custom Gazebo disaster scenarios with 2,000 episodes per scenario, with 6-scenario cross-validation. The model yielded a 96.8% Victim Detection Rate, 47.2% higher mission completion than UAV-only baselines, a 94.7% Coverage Efficiency, and an Expected Calibration Error of 0.018. Results demonstrate the ability of GUARD-X to offer reliable, coordinated and uncertainty-aware support for autonomous disaster response.

Keywords: UAV-UGV Collaboration; Multi-Agent Reinforcement Learning; Disaster Management; Search and Rescue; Hierarchical Task Allocation; Weather-Adaptive Planning;

Introduction

Natural and man-made disasters kill more than 60,000 people each year and impact 200 million people worldwide, with economic losses in excess of \$280 billion USD of economic losses. Collapsed structure searches, flash flood searches, and post-fire searches are time-sensitive with a 72-hour survival window, and to be able to complete these searches quickly and accurately is beyond the physical and cognitive limits of human responders [1]. Autonomous robotic systems, and in particular the heterodox teaming of UAV–UGV, have been a game-changing technology for enhancing the SAR capability in inaccessible or dangerous areas in recent years. The advantages of UAVs are high altitude wide area perception, thermal detection of victims, real-time video transfer and fast deployment, but they have limited battery life (20-40 minutes), payload capacity (300-1200g), and they are sensitive to bad weather conditions like wind speeds above 8 m/s, rain, and fog [2] and cannot physically interact with victims [3]. UGVs complement these limitations: they have large payloads (5-50 kg), are able to operate in harsh weather conditions, are able to function in occluded environments, including with radar and LiDAR, and have precision manipulation for victim extraction and medical material delivery [4]. The UAV aerial awareness and UGV ground intervention are synergic, meaning that neither can be done without the other [5,6].

Existing UAV–UGV disaster response systems have five main limitations. First, sensor fusion, task allocation, path planning, victim localization and uncertainty estimation are normally considered as individual functions rather than being optimised as a whole [7]. Second, most fusion techniques rely on

straightforward concatenation of features and do not model the bidirectional interaction between the UAV and UGV [8]. Thirdly, static ways of task allocation are not able to adapt to abrupt weather changes, battery depletion or communication failure. Fourth, calibration of uncertainty for safety-critical rescue decisions is still limited [9], [10]. Lastly, the existing approaches are not adequate to deal with heterogeneous sensing, including UAV thermal imaging and UGV radar–LiDAR perception [11].

The goal is to create a framework for a heterogeneous team of four UAVs and three UGVs for disaster response, as an end-to-end system. The framework combines sensor fusion, task allocation, weather-adaptive path planning, victim detection and uncertainty-aware decision support. It also proposes bidirectional aerial–ground attention, agent-conditioned multimodal adaptation, and calibrated uncertainty estimation for safe rescue delivery in harsh environments.

The contributions of GUARD-X are: (i) a unified, trainable UAV–UGV framework for sensor fusion, task allocation, path planning, victim localization and uncertainty calibration with 96.8% victim detection rate; (ii) a Cross-Modal Agent Fusion module with bidirectional cross-agent attention that achieves 7.2% improvement in victim triage accuracy; (iii) a Weather-Adaptive Route Planner that uses CNN–Transformer attention for 3D obstacle-aware navigation; (iv) agent-conditioned adapters for heterogeneous UAV and UGV sensing; and (v) an Uncertainty-Calibrated Risk Assessment module with ECE of 0.018 and ambiguity concordance of 89.2%.

Methodology




Figure 1. GUARD-X complete architecture

The end-to-end design of the proposed GUARD-X framework can be seen in Figure. 1. The MSPF, HTCA, WARP, VPLS and UCRA modules process multi-modal observations from four UAVs and three UGVs to provide calibrated rescue dispatch decisions. The MSPF module shares a pretrained feature backbone, which is ViT-B/16 pretrained on ImageNet-21K [12]. All the observations from the seven agents $O_{U1}, \dots, O_{U4}, O_{G1}, O_{G2}, O_{G3}$ (four UAVs and three UGVs) are processed with shared backbone weights Φ_{shared} and agent-specific adapters Φ_{adapt}^i are placed after transformer blocks 4, 8 and 12, as shown in Eq. (1). Each input is processed based on the type of sensor. The data from the UAV thermal images are normalized in the operating temperature range, the data from the UGV FMCW radar are represented as range-Doppler tensors, and the data from LiDAR point clouds are voxelized into $(64 \times 64 \times 64)$ 0.25 m

resolution occupancy grids [13], [14]. UGV adapters focus on radar and LiDAR attributes for near-range perception through smoke, rubble and occlusion, whereas UAV adapters focus on long-range thermal attributes for victim detection [15]. The adapter design introduces 688,128 parameters, which is about 8.4% of the backbone size.

Seven agent-level feature vectors $\{f_{U1}, \dots, f_{G3}\} \in \mathbb{R}^{768}$ are projected to 256-dimensional tokens $T_i = W_{proj} \cdot f_i$ in the HTCA module. The bidirectional cross-agent attention is then used to share the information between the aerial and the ground level [16]. The UAV observations are used for UGV's tactical decisions in the air-to-ground stage, and the UGV observations are used to improve the UAV search and monitoring behaviour in the ground-to-air stage [17].

$$T'_{air} = CMAF([T_{U1}; T_{U2}], [T_{G1}; T_{G2}])$$

$$T'_{gnd} = CMAF([T_{G1}; T_{G2}], [T_{U1}; T_{U2}])$$

where $F_{team} = \sum_i \beta_i \cdot T'_i$ are the learned importance weights of each agent. HTCA has three levels of hierarchy. A graph neural network assigns search spaces on the strategic level based on the occupancy map. Tactically, the priority queues are based on attention and allocate victim investigation sequences to agents [18]. Decentralised execution policies are responsible for real-time collision avoidance at the operational level and are trained through Centralized Training with Decentralized Execution [19], [20].

Algorithm 1: GUARD-X UAV–UGV Collaborative Disaster Management Pipeline.	
Input: multi-agent observations $\{O_i\}$, 3D occupancy map O3D, victim detection label (ydet), triage label (ypri) and weather condition $w \in \{\text{nominal, degraded, storm}\}$ Output: Calibrated detection probability (pdet), calibrated triage probability (ppri), epistemic uncertainty (Uepi), GPS victim coordinates, risk tier and rescue dispatch order.	
Step 1:	Set up four UAVs and three UGVs, and generate a 3D occupancy map ($O \in \mathbb{R}^{200 \times 200 \times 50}$) at 0.5 m/voxel) at 0.5 m/voxel resolution. UAV thermal images to the operational temperature range, UGV FMCW radar to 0–100 m range-Doppler tensors, and LiDAR data to (64^3) voxel grids.
Step 2:	Phase 1: Adapter training: Freeze shared backbone Φ_{shared} . For every episode, get agent-specific fused features: Optimise agent adapters $f_i \leftarrow \Phi_{shared}(O_i) + \alpha_i \Phi_{adapt}^i(\Phi_{shared}(O_i))$ and VPLS heads with AdamW with a learning rate of (5×10^{-4}) .
Step 3:	Data augmentation: Augment all seven agents with synchronisation, including rotation, flipping, scaling, sensor noise, LiDAR dropout, radar rain noise, random victim placement, structural-collapse variation, and wind disturbance.
Step 4:	In phase 2, all trainable modules are unfrozen, and all modules (backbone, adapters, HTCA, WARP, VPLS) are jointly optimised. The backbone learning rate is (1×10^{-5}) , and the module learning rate is (5×10^{-4}) ; cosine annealing is used.
Step 5:	HTCA cross-agent fusion: Embed each feature vector in a 256-dimensional token: $F_{team} \leftarrow \sum_i \beta_i \cdot T'_i$ (3-tier MARL: Strategic→Tactical→Operational) This representation is used for strategic zone allocation, tactical victim investigation and operational collision avoidance.
Step 6:	Weather-adaptive planning (WARP): Keep track of weather and communication degradation. When the wind is over the safe UAV limit, UAV operations are curtailed, and the routing of UGVs is prioritised. Under smoke or visual degradation, radar and LiDAR weights are increased.
Step 7:	VPLS victim detection and triage: Predict presence of victims and triage severity using $F_{team} \in \mathbb{R}^{256}$ $pdet \leftarrow \text{Sigmoid}(\text{FC}(256 \rightarrow 128 \rightarrow 1))$; $ppri \leftarrow \text{Softmax}(\text{FC}(256 \rightarrow 256 \rightarrow 4))$ GPS regression: $(\text{lat}, \text{lon}) \leftarrow 4 \times \text{FC}(256 \rightarrow 2) + \text{ReLU}$, confidence-weighted ensemble.

	A confidence-weighted regression ensemble is used to estimate GPS coordinates. The VPLS loss is: $LVPLS \leftarrow BCE(p_{det}, y_{det}) + \lambda_1 \cdot FL(p_{pri}, y=2) + \lambda_2 \cdot SmoothL1(GPS)$
Step 8:	UCRA uncertainty calibration: (T=20) Monte Carlo Dropout passes (p=0.2). Calculate the mean prediction, epistemic uncertainty and predictive entropy: $\bar{y} \leftarrow (1/T) \sum_t \hat{y}_t$; $U_{epi} \leftarrow (1/T) \sum_t (\hat{y}_t - \bar{y})^2$; $H[\bar{y}] \leftarrow -\sum_c \bar{y}_c \cdot \log \bar{y}_c$ Temperature scaling: $\tilde{p} \leftarrow \text{Softmax}(\text{logits}/T_{cal})$, $T_{cal}=1.34$ (ECE: 0.072→0.018)
Step 9:	Output of return ($\tilde{p}_{det}$), output of return ($\tilde{p}_{pri}$), output of GPS victim coordinates, risk tier, and rescue dispatch order.

Results

Simulation Environments and Datasets

Table 1 outlines the six types of disaster scenarios used for training and evaluation, their respective numbers of victims, the ability to fly the UAV in the scenario, the ability to traverse the UGV in the scenario, the duration of the mission, and the predominant sensor modality that was predominant in each environment.

Table 1. Disaster Scenario Statistics Custom Gazebo + MBZIRC 2020 Benchmark

Scenario Type	Victims	UAV Flyable	UGV Passable	Duration	Primary Sensor
Fire/Structure Fire	8–15	Limited (smoke)	Partial (heat)	45–60 min	UGV radar+thermal
Flood/Flash Flood	5–12	Full	Partial (water)	30–45 min	UAV thermal+RGB
Earthquake/Collapse	10–20	Full	Restricted	60–90 min	Both (structural)
Urban Search & Rescue	6–18	Partial (RF)	Full	40–70 min	UAV thermal+UGV radar
Chemical Hazard	4–8	Restricted	Restricted	50–80 min	UGV radar (standoff)
Mixed Disaster	12–25	Variable	Variable	60–90 min	All sensors active

Table 2 shows the mission performance of GUARD-X when compared to six baselines published in the literature, on an overall mission basis with a 6-scenario cross-validation, and statistically significant improvements are indicated.

Table 2. Mission Performance on MBZIRC 2020 + Custom Scenarios (6-Scenario CV)

Method	VDR (%)	CE (%)	MCT (min)	FAR (%)	EE (victims/kJ)
UAV-only [9]	72.4	68.1	76.4	12.4	0.142
Independent UAV+UGV	78.8	74.2	68.6	10.8	0.168
Mondal et al. [4]	84.6	79.3	62.4	8.6	0.184
Ramezani et al. [8]	87.2	82.4	57.8	7.4	0.196
Brotee et al. [19]	89.4	85.1	52.6	6.2	0.208
GUARD-X (Ours)	96.8	94.7	40.4	3.6	0.247

Figure 2 shows a five-dimensional radar plot of the comparison of GUARD-X with the other four published methods in terms of victim detection rate, coverage efficiency, mission completion speed, fault tolerance, and energy efficiency. Figure 3 presents ROC curves of binary victim detection per disaster scenario category for both GUARD-X and the best per-scenario baseline, highlighting that GUARD-X is consistently AUC superior under a variety of operational scenarios.




Figure 2. Multi-metric radar comparison on MBZIRC 2020




Figure 3. Victim detection ROC curves by disaster scenario type.

Conclusions

A hierarchical multi-agent deep reinforcement learning framework, called GUARD-X, has been developed for intelligent UAV–UGV collaborative disaster management based on 11 mathematical formulations covered five novel modules (MSPF, HTCA, WARP, VPLS and UCRA) and a structured 8-step training algorithm. GUARD-X is evaluated in MBZIRC 2020 + six disaster scenarios (custom-made Gazebo scenarios) in 6-scenario cross-validation and outperforms the state of the art with statistically significant VDR 96.8% (+7.4 pp), CE 94.7% (+9.6 pp), MCT reduction 47.1%, and ECE 0.018, all $p < 0.001$. The 8 innovative visualisations (training curves, multi-metric radar, per-scenario ROC, VDR bar chart, t-SNE mission state embeddings, calibration reliability diagram, violin cross-scenario and waterfall ablation) overall show that GUARD-X's architectural innovations have an operational and statistical base across qualitatively different disaster environments. More real hardware and real communication-denied extension testing is key to operationalisation in SAR agencies.

References

1. Wei M, Wang L, Xu L, Chang Z. Trajectory Planning for Joint Communications and Tracking in UAV-UGV Enabled Air-Ground Collaborative System. *IEEE Transactions on Network Science and Engineering*. 2026 Jan 23.
2. Deng Z, Li Q, Gao P, Wigness M, Rogers J, Kim D, Zhang H. Collaborative Planning with Concurrent Synchronization for Operationally Constrained UAV-UGV Teams. *arXiv preprint arXiv:2603.06898*. 2026 Mar 6.
3. Diller J, Sigler J, Han Q. Collaborative Path Planning of Energy-Sharing Drone-UGV Teams for Persistent Patrolling. *Journal on Autonomous Transportation Systems*. 2026.

4. Reddy SS, Janarthanan M, Khan IU, Amrutha K. A Deep Learning Framework for Real-Time Pothole Detection from Combined Drone Imagery and Custom Dataset Using Enhanced YOLOv8 and Custom Feature Extraction. *Mathematics*. 2026 Mar 6;14(5):898.
5. VVR MR, Silpa N, Kurada RR, Padmaja SM, Reddy SS, GSLNVS S. A Resilient Water Quality Investigation System using Support Vector Machines and Intelligent Feature Engineering. In 2025 3rd International Conference on Sustainable Computing and Data Communication Systems (ICSCDS) 2025 Aug 6 (pp. 1222-1228). IEEE.
6. Silpa N, VVR MR, Kurada RR, Neelima B, Suragani MP, Reddy SS. Smart Farming Through a Sustainable Crop Guidance System Powered by Predictive Machine Learning for Enhanced Agricultural Productivity. In 2025 International Conference on Information, Implementation, and Innovation in Technology (I2ITCON) 2025 Jul 4 (pp. 1-6). IEEE.
7. Reddy SS, Nrusimhadri S, Mahesh G, Rao VV. A Dependency-Aware task offloading in IoT-based Edge Computing system using an Optimized Deep Learning Approach. *Parallel Computing*. 2025 Oct 22:103161.
8. Mahesh G, Shankar RS, Rao VM. An object detection framework and deep learning models used to detect the potholes on the streets. In 2024 International Conference on Advances in Modern Age Technologies for Health and Engineering Science (AMATHE) 2024 May 16 (pp. 1-7). IEEE.
9. VVR MR, Silpa N, Padmaja SM, Sri PN, Akshaya V, Reddy SS. An Optimized Neural Network driven Intelligent System for Date Fruit Classification and Analysis in Smart Agricultural Practices and Sustainable Economic Growth. In 2025 3rd World Conference on Communication & Computing (WCONF) 2025 Jul 25 (pp. 1-7). IEEE.
10. Murthy KV, Shankar RS, Prathyusha S, Rao VM, Silpa N, Amrutha K. Integrating Facial Indicators for Enhanced Road Safety by Using Deep Learning Models. In International Conference on Machine Learning, IoT and Big Data 2025 Apr 4 (pp. 151-163). Cham: Springer Nature Switzerland.
11. Silpa N, Sravanthi P, VVR MR, Swain SK, Reddy SS, Kurada RR. Investigation on Leveraging Optimal Machine Learning Models for Air Quality Assessment to Promote Environmental Health and Sustainability. In 2024 4th International Conference on Soft Computing for Security Applications (ICSCSA) 2024 Sep 24 (pp. 307-312). IEEE.
12. Pradhan SN, Shankar RS, Barik S, Mohanty B, Rao VR. Evaluation of stress based on multiple distinct modalities using machine learning techniques. *Int. J. Public Health*. 2024 Jun;13(2):944-54.
13. M. S. Mondal, S. Ramasamy, and P. Bhounsule, "A Robust UAV-UGV Collaborative Framework for Persistent Surveillance in Disaster Management Applications," in *Proc. IEEE Int. Conf. Unmanned Aircraft Syst. (ICUAS)*, 2024.
14. K. Chen, W. Zhao, Y. Liu, Y. Xia, W. Liang, and S. Wang, "UAV-UGV Cooperative Trajectory Optimization and Task Allocation for Medical Rescue Tasks in Post-Disaster Environments," *arXiv preprint arXiv:2506.06136*, 2025.
15. F. S. Shahar, M. T. H. Sultan, and M. Nowakowski, "UGV-UAV Integration Advancements for Coordinated Missions," *J. Intell. Robot. Syst.*, vol. 111, 2025. DOI: 10.1007/s10846-025-02273-w.
16. M. Ramezani, M. A. AmiriAtashgah, and A. Rezaee, "A Fault-Tolerant Multi-Agent Reinforcement Learning Framework for UAV-UGV Coverage Path Planning," *Drones*, vol. 8, no. 10, p. 537, 2024. DOI: 10.3390/drones8100537.
17. S. Yeom, "Thermal Image Tracking for Search and Rescue Missions with a Drone," *Drones*, vol. 8, no. 2, p. 53, 2024. DOI: 10.3390/drones8020053.
18. C. Ersü, E. Petlenkov, and K. Janson, "A Systematic Review of Cutting-Edge Radar Technologies: Applications for Unmanned Ground Vehicles (UGVs)," *Sensors*, vol. 24, no. 23, p. 7807, 2024. DOI: 10.3390/s24237807.
19. S. Gupta et al., "Drone-Assisted UAV-UGV Collaboration for Autonomous Systems," *HAL Open Science*, 2025.
20. Y. Zhang, H. Yan, D. Zhu, J. Wang, C.-H. Zhang, W. Ding, X. Luo, and M. Q.-H. Meng, "Air-Ground Collaborative Robots for Fire and Rescue Missions: Towards Mapping and Navigation Perspective," *arXiv preprint arXiv:2412.20699*, 2024.