Federated DDQN for Edge-based Disaster Response: Leveraging the Extensible Front-line Augmented Communication Exchanger (X-FACE)

Babatunde Ojetunde†, and Toshikazu Sakano†
†*Wave Engineering Laboratories, Advanced Telecommunications Research Institute International*
Seika, Soraku, Kyoto, Japan
{ojetunde, t.sakano }@atr.jp

***Abstract*—The adoption of artificial intelligence (AI) in disaster countermeasures is transforming emergency response, yet real-world frontline deployment remains limited by unstable connectivity, scarce computing resources, power constraints, and non-expert users. This paper introduces a Federated Deep Double Q-Network (DDQN) reinforcement learning model deployed on the Extensible Front-line Augmented Communication Exchanger (X-FACE), an edge-based portable device. The proposed model autonomously maps disaster-affected zones, enabling agents to learn efficient search policies and support first-responder decision-making even in fully offline or resource-limited conditions. By continuously optimizing search routes and task allocation, the framework accelerates the detection and rescue of missing survivors. Experimental results show that the proposed method achieves high detection performance (AUDC up to 0.82), enables early survivor discovery within roughly 3 hours, and recovers more survivors in shorter missions, demonstrating practical effectiveness for rapid and robust search and rescue operations.**

***Index Terms*—Artificial Intelligent, DDQN, portable edge server, disaster search and rescue,IoT, wearable device, communication, coordination activities**

I. Introduction

The frequency and severity of natural disasters have intensified in recent years, and in most of these events, effective communication plays a crucial role in facilitating timely and coordinated search and rescue efforts. However, in many cases of large-scale disasters, telecom services and Internet connectivity are disrupted due to infrastructure damage or traffic congestion. In addition, first responders (e.g., firefighters) face significant challenges in carrying out their duties of searching, rescuing and ensuring the safety of survivors in disaster-affected areas.

To address these challenges we have proposed a Front-Line Operation System (FLOS) [1], later renamed to the Extensible Frontline Augmented Communication Exchanger (X-FACE, referred to as X-FACE hereafter). X-FACE enables frontline operators such as firefighters, medical personnel, and police officers to efficiently collect and share information at disaster sites. The system incorporates edge-based AI functions that provide features such as speech recognition, intelligent image processing, and automatic analysis, and extracting information even under Internet disruption. To further improve disaster response operations, particularly in search and rescue, and to ensure that the challenges faced by first responders are alleviated, an edge-based federated Double Deep Q-Network (DDQN) [2] is introduced as an important component.

This paper introduces a novel disaster response architecture that integrates a resource-efficient Federated Dueling Double Deep Q-Network (Federated DDQN) deployed on the X-FACE. The proposed framework empowers multiple edge-based agents, such as first responder teams, drones, and robots, to autonomously learn, adapt, and coordinate search actions using local sensing and communication. The proposed Federated DDQN combine Federated aggregation, adaptive quantization, and advanced reinforcement learning, to delivers an operationally practical, scalable, and field-ready AI solution. Unlike traditional centralized models, the proposed method distributes intelligence across multiple agents, creating a collective intelligence network that strengthens search and rescue policy. The performance and advantages over conventional search paradigms and edge model reduction techniques are validated in realistic multi-agent search and rescue evaluation. The goal of the proposed method is to facilitate improved search and rescue operations with minimal constraints.

The rest of this paper is organized as follows. Section II, presents the related works on the use of AI in disaster response. Section III describes the proposed federated DDQN. In Section IV, we describe the performance evaluation of the proposed method, and Section V concludes the entire paper.

II. Related Work

The research related to the use of AI in disaster response operation has been gaining increasing focus in recent times. Merkle et al. [3] focuses on utilizing drone-based data and deep learning methods to automate situation assessment and aid delivery in disaster relief efforts.

Samikwa et al. [4] presented a system that combines IoT, Artificial Neural Networks (ANN), and edge computing for short-term flood prediction. The proposed system used LSTM to predict flood water levels ahead of time by monitoring rainfall and water level sensor data in real-time. Stateczny et al. [5] also introduced a novel deep hybrid model for flood prediction that combines Convolutional Neural Network (CNN) and Residual Network (ResNet) classifiers. It focuses on flood detection using remote sensing satellite images and incorporates preprocessing techniques like median filtering and segmentation. Similar to this, Alsumayt et al. [6] introduced a method for flood detection in Saudi Arabia using drones equipped with a Flood Detection Secure System (FDSS). They integrates deep active learning based classification models and blockchain-based federated learning with partially homomorphic encryption (PHE).

Conversely, Abid et al. [7] reviewed recent studies on the applications of artificial intelligence (AI) in disaster risk reduction, disaster preparedness, and disaster response. The review offers a comprehensive overview of the current state and trends of AI research in disaster management. Andreassen et al. [8] discussed the challenges faced by emergency response coordination in complex environments, such as the Arctic, where limited resources, unpredictable conditions, and technical limitations hinder information sharing among response units.

While efforts are being made to use AI in disaster countermeasures, there are various challenges such as data reliability and quality, AI integration and compatibility, and resource allocation and accessibility. As part of our contribution, our proposed method is implemented with these challenges in mind and can help reduce some of these constraints. Furthermore, our Federated DDQN differs from state-of-the-art methods by introducing an innovative architecture that combines a resource-efficient Dueling Double Deep Q-Network (DDQN) with a Federated Deep Reinforcement Learning (FDRL) framework, enabling hyperparameter optimization and improving overall search and rescue performance.

III. Proposed Federated DDQN for Disaster Search and Rescue

In order to improve efficiency in search and rescue operations under disaster situations where communication and ICT services are disrupted, it is important to have advanced ICT tools and mechanisms that can support multi-responder missions, where efficient coordination is essential and decisions must be made under partial observability. Therefore, in this paper, we propose a Federated DDQN model deployed on X-FACE (i.e., a portable system embedded with edge-based AI functions).

To achieve this, in the aftermath of a large-scale disaster, a first responder team (e.g., a team leader and four members) is deployed to the affected area for a search and rescue mission. The team leader carries X-FACE (i.e., a wearable device) when searching for survivors in the disaster area. Each first responder operates within the network coverage radius of the X-FACE core.

Unlike traditional approaches, where there is no means of knowing which locations have already been searched by each first responder, X-FACE provides efficient search coordination. This is made possible through its embedded AI functions. The search area of each first responder is automatically analyzed and shared with all team members via the X-FACE core. In addition, X-FACE ensures that all voice-based communication can be automatically transcribed into text using its speech-to-text engine. This allows first responders to conduct hands-free searches without the need to manually input information on their devices, thereby improving the processing time required to carry out their activities.

Each first responder's device acts as an agent in a multi-agent reinforcement learning system. These agents learn optimal search patterns and strategies over time. The X-FACE core (miniPC) serves as a coordinating hub, dynamically allocating search areas to each agent based on the current location of each responder, areas already searched, the probability distribution of survivor locations (updated in real time), and individual agent performance metrics.

Therefore the following subsections outline the complete Federated DDQN pipeline, from local agent design to model formation, federated aggregation, portability, and final training and deployment.

A. Dueling DDQN for Local Edge Agents

Building on the coordination capabilities provided by X-FACE, in this section we explain the underlying reinforcement learning model that enables agents to learn and adapt search strategies in real time. Each frontline device operates as a local edge agent powered by a Federated DDQN.

In our proposed model, each first responder's device functions as an intelligent agent embedded within a cooperative search and rescue mission. In this setting, we adopt a Dueling Double Deep Q-Network (DDQN), which decomposes the value estimation into separate streams: one for the value function $V(s)$ and another for the advantage function $A(s, a)$. Here, the term "streams" refers to the two internal neural network branches that separately compute the state-value $V(s)$ and the action-advantage $A(s, a)$ before they are combined to form the final Q-value. For state s and action a, the Q-value is derived as in Equation 1, where $\theta = \{\alpha, \beta\}$ represents the set of trainable weights, with α corresponding to the parameters of the advantage stream and β corresponding to the parameters of the value stream. The function $V(s; \beta)$ estimates the value of being in state s, while $A(s, a; \alpha)$ estimates the relative advantage of taking action a in state s. The subtraction of the mean advantage term $\frac{1}{|A|}\sum_{a'} A(s, a'; \alpha)$ ensures that the advantage function is centered, preventing identifiability issues between the value and advantage streams. Here, A denotes the set of possible actions available to the agent in the current situation, such as movement directions, communication actions, or assisting survivors.

$$Q(s, a; \theta) = V(s; \beta) + \left(A(s, a; \alpha) - \frac{1}{|A|} \sum_{a'} A(s, a'; \alpha) \right) \tag{1}$$

The design of our proposal enables sub-256 KB model storage and energy-efficient execution, ensuring that the Federated DDQN remains practical for real-time disaster response operations.

B. Model Formation for Search and Rescue

Based on the local DDQN architecture for edge agents formed in III-A, this section describe how the model is formed for search and rescue missions, detailing the state representations, reward design, and action selection process.

Each agent periodically senses its local environment—mapping features such as survivor likelihood heatmaps, map accessibility, teammate locations, and input from nearby X-FACE units. The agent then selects actions to maximize long-term mission reward, which is computed as a weighted combination of survivor finds, time penalties, battery conservation, and communication success. The reward guides the agent by assigning positive values to actions that lead to survivor discovery and efficient coverage, and negative values to actions that waste time or revisit explored areas, enabling the DDQN to learn search behaviors that maximize long-term mission success.

In order for the model to adopt a robust policy that can be deployed towards disaster search and rescue, the policy learning proceeds online, with the DDQN parameters updated after each episode using the double-network target update as shown in Equation 2. Here, α denotes the learning rate, θ^{-} represents the target network parameters, and γ is the discount factor controlling reward foresight. In addition, Q_θ refers to the online Q-network, while Q_{θ^-} denotes the target Q-network. These local updates are subsequently aggregated across responders through X-FACE, forming the basis of the Federated DDQN for coordinated search and rescue.

$$\theta \leftarrow \theta + \alpha \Big[r_t + \gamma\, Q_{\theta^-}\Big(s_{t+1}, \arg\max_{a'} Q_\theta(s_{t+1}, a')\Big) - Q_\theta(s_t, a_t)\Big] \nabla_\theta Q_\theta(s_t, a_t) \quad (2)$$

Figure 1 shows the operational flow of our Federated DDQN model deployed on X-FACE. The area data from the disaster grid is sensed and uploaded to the X-FACE core, where local agents perform scans and receive coordinated assignments. Each agent trains a local DDQN model, which is periodically aggregated via a federated learning protocol. The updated global policy is redistributed to agents, improving survivor detection and route guidance. This cycle enables adaptive, distributed learning under resource constraints, as validated in our evaluation.

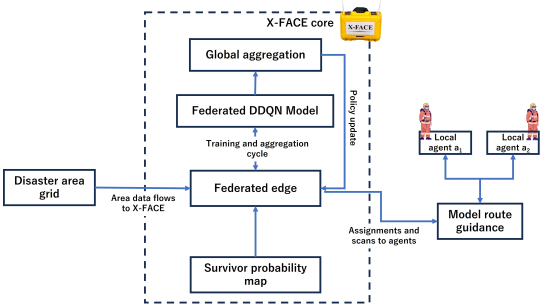

Fig. 1. Federated DDQN architecture.

C. Federated DDQN Protocol

Since search and rescue missions demand cooperation without dependence on centralized servers, it is necessary to adopt federated deep reinforcement learning as part of our proposed system. Each X-FACE periodically transmits its updated (locally quantized) DDQN parameters θ_i to a secure FDRL aggregation node. The FDRL server updates the global model as a reliability-weighted mean, as shown in Equation 3. Here, θ_i denotes the locally trained DDQN parameters of agent i, θ_{global} represents the aggregated global parameters, w_i is the reliability weight assigned to agent i, and N is the total number of participating agents. The weights w_i reflect real-time communication success, battery level, and local episode return, thereby promoting contributions from more reliable and effective units. After aggregation, the improved global parameters θ_{global} are redistributed, and agents resume local training with the updated model, ensuring adaptation to both global and local trends.

$$\theta_{\text{global}} = \frac{\sum_{i=1}^{N} w_i \cdot \theta_i}{\sum_{i=1}^{N} w_i} \quad (3)$$

This protocol allows every agent to benefit from collective learning, even when direct peer-to-peer synchronization is disrupted by intermittent wireless or environmental conditions. Unlike conventional federated averaging or static parameter sharing, the proposed Federated DDQN protocol integrates operational reliability and real-time performance directly into the aggregation process.

D. Federated DDQN Model Portability

To achieve the portability of our proposed Federated DDQN, we adopt Adaptive Action-Set Quantization (AASQ) to compress the model and enable smooth execution in resource-constrained environments. Typical quantization applies uniform (e.g., 8-bit) compression to all network weights, which can penalize accuracy for actions or features that are critical under local mission constraints. Our AASQ mechanism introduces a dynamic, action-aware scheme that selectively allocates precision to network branches associated with high-variance, and important actions. Formally, after each learning round, the bitwidth b_j for layer j is set as in Equation 4:

$$b_j = b_{min} + (b_{max} - b_{min}) \cdot \sigma(Q^j)^p \quad (4)$$

Here, $\sigma(Q^j)$ is the normalized standard deviation of Q-values at the output neurons of layer j, computed over a sliding observation window during search and rescue operations. The parameters b_{min} and b_{max} define the adaptive bitwidth range (e.g., 4–8 bits), while p modulates sensitivity. For branches tied to critical and uncertain actions (high $\sigma(Q^j)$), the mechanism preserves higher precision; for stable, routine actions, it reduces bitwidth, minimizing model footprint and memory operations. By analyzing reward and action-value variance in real time, AASQ customizes quantization to the operational context rather than applying a one-size-fits-all policy.

E. Training and Deployment

Training of the proposed model begins with each agent's DDQN using simulated, stochastic disaster scenarios, followed

TABLE I
SIMULATION PARAMETERS

Parameter	Value
Disaster area sizes	$400\,\mathrm{m} \times 400\,\mathrm{m}$ up to $1\,\mathrm{km} \times 1\,\mathrm{km}$
Grid cell size	$20\,\mathrm{m}$
Training grid sizes	20×20, 30×30, 40×40
Evaluation grid sizes	10×10, 15×15, 25×25, 35×35
Number of agents	4–8 (scenario-dependent)
Number of survivors	15–35 (scenario-dependent)
Detection probability	0.6–0.85
Scan requirement	2–4 scans
Scan range	2–3 grid cells (Manhattan distance)
Obstacle ratio	0.2–0.3
Search time (evaluation)	6–24 hours
Agent movement speed	$1.0\,\mathrm{m/s}$
Replay buffer size	10^5 transitions
Batch size	128
Learning rate	0.0005
Discount factor (γ)	0.95
Model aggregation interval	Every 1000 agent steps

by federated learning rounds in which compressed parameter updates is shared with X-FACE for aggregation. The AASQ is applied after every local update round, ensuring parameters are stored and communicated efficiently. The final quantized federated model is then deployed to X-FACE cores, enabling real-time inference and robust responder coordination.

The proposed model hyperparameters are selected to reflect edge deployment realities. For example, the learning rate is set to $\alpha = 0.0005$, the discount factor to $\gamma = 0.95$, and the replay buffer includes 10^5 transitions, with a batch size of 64 and a synchronization interval of every 1000 agent steps. For quantization, the bitwidth bounds are set to $b_{min} = 4$ and $b_{max} = 8$, with the adaptation parameter $\rho = 0.5$, ensuring a balance between efficiency and expressivity for the deployed models.

IV. PERFORMANCE EVALUATION

In this section, the performance evaluation of the proposed federated DDQN for disaster search and rescue was conducted through a structured model training and evaluation process implemented in Python programming. The model training scenarios were designed to expose agents to varying grid sizes, survivor densities, obstacle ratios, and detection probabilities, that reflects realistic search conditions and agents limitations (here agents are physical first responders/drone/sensors), while the evaluation scenarios are designed to extend these conditions to longer search times and stricter scan requirements. Each scenario was designed to simulate a disaster-affected area, where missing survivors are initially placed using a uniform random distribution across a defined grid. The grid represents a physical area which were discretized into grids of different sizes.

Each first responder is modeled as a mobile agent with access to the X-FACE core, operating under a random walk mobility model with a strict non-revisit constraint. Agents avoid previously visited cells and those visited by other team members, promoting spatial coverage and reducing redundancy. The movement speed was set to 1.0 m/s, and agents perform active scans for survivors within their scan range. The scanning process incorporates probabilistic detection and multi-scan confirmation, reflecting real-world limitations in thermal imaging and acoustic sensing. Survivors are only confirmed after a required number of scans, and detection probabilities vary across scenarios to simulate sensor degradation and environmental interference.

Hyperparameters for training were selected to reflect practical edge settings. The federated model aggregation every 1000 agent steps. The evaluation phase was conducted using both CPU and GPU systems, with scenario files specifying search durations in hours, detection probabilities, scan requirements, and obstacle ratios. The summary of the key parameters used in both training and evaluation is shown in Table I.

A. Model training results

To determine the usability of the proposed model in improving disaster search and rescue operation, first we train the model using a multi-grid approach. Figure 2 shows the model training metrics.

1) Training Loss:

The training loss profile of the proposed federated DDQN is shown in Fig. 2(a). As shown in the results, the loss fluctuated significantly during the early stages of training, particularly between update steps 1000 and 4000. This instability is expected due to the initial exploration phase, where agents interact with diverse scenarios and the replay buffer is still being populated. The sharp spikes in loss during this phase reflect the high variance in Q-value targets, especially when agents encounter rare survivor configurations or high obstacle densities.

After update step 4000, a noticeable decline in loss values was observed, indicating that the model began to stabilize and converge. The reduction in loss suggests that the Q-network was able to approximate the expected returns more consistently across episodes. The final segment of the curve shows a smoother trajectory, confirming that the federated updates and target network synchronization were effective in reducing gradient noise and improving convergence.

2) Temporal Difference (TD) Error:

Furthermore, we measure the TD error, which measures the discrepancy between predicted Q-values and target returns, as shown in Fig. 2(b). The curve reveals a dynamic learning process, with TD error rising steadily during the initial updates and peaking around step 3000. This increase corresponds to the agent encountering more complex survivor placements and obstacle configurations, particularly in the larger grid scenarios. The elevated TD error during this phase reflects the challenge of generalizing across multiple environments with varying scan requirements and detection probabilities.

Following the peak, the TD error began to decline, indicating that the agent was able to reduce prediction inconsistencies and improve its value estimation. These fluctuations are not detrimental; rather, they highlight the system's ability to adapt to diverse conditions without overfitting to a single environment. The overall trend confirms that the agent was learning

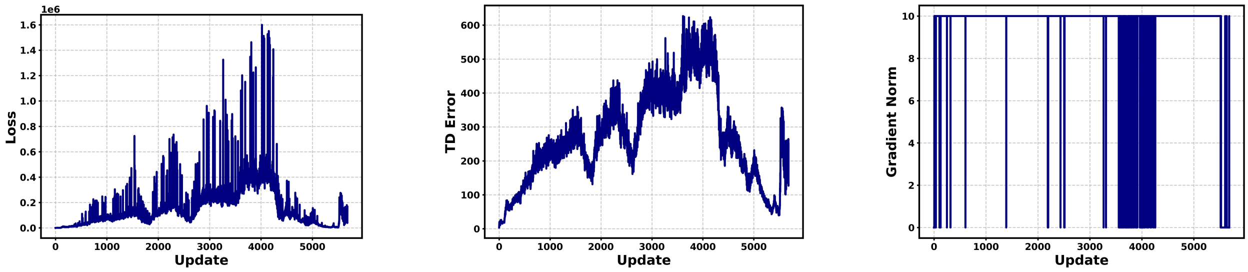

(a) Training loss. (b) Temporal difference error. (c) Gradient Normal.

Fig. 2. Model training.

to align its Q-values with actual returns, a critical aspect of stable policy improvement.

3) Gradient Norm:

We consider the gradient norm behavior, which is used to manage and balanced magnitude of parameter updates and the stability of the optimization process during model training. As shown in Fig. 2(c), the gradient norm remained close to the clipping threshold of 10 for a substantial portion of the training, particularly between update steps 1000 and 4000. This indicates that the model was undergoing aggressive updates, likely due to high TD error and volatile Q-value estimates during early training.

The frequent spikes to the maximum value suggest that the agent was encountering steep gradients, which were clipped to prevent instability. This behavior is consistent with the use of federated aggregation and scenario cycling, where abrupt changes in environment dynamics can lead to sharp shifts in the loss landscape. The reduced gradient activity in the later stages confirms that the agent was refining its policy with smaller, more precise updates, a sign of convergence and improved generalization.

4) Survivors Found and Coverage Ratio:

Lastly, the task-level performance of the agent was evaluated using the number of survivors found per episode and the coverage ratio, as shown in Fig. 3. The dual-axis plot presents both metrics over 100 training episodes, providing a clear view of the agent's search effectiveness and spatial exploration.

The results confirm that initially, the number of survivors found per episode showed a rising trend, indicating that the agent was learning to prioritize survivor detection. However, beyond episode 40, a gradual decline was observed. This reduction may be attributed to the agent's shift toward exploitation, where it favors known high-reward paths and reduces exploratory spread. The coverage ratio followed a similar trajectory, with early episodes showing broad exploration and later episodes exhibiting more focused movement. This behavior is consistent with the epsilon decay schedule and the influence of revisit penalties and proximity shaping in the reward function. Despite the decline, the agent maintained a reasonable balance between detection and coverage, suggesting that it was able to adapt its strategy based on scenario complexity.

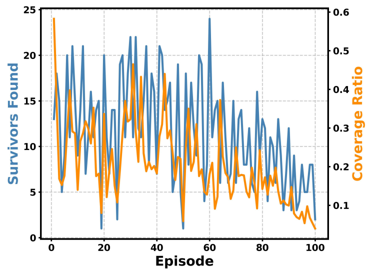

Fig. 3. Survivors found vs search coverage.

B. Model evaluation

In order to confirm the effectiveness of our proposed federated DDQN model, we evaluate its performance when deployed in a search and rescue operation. For this, we used four different scenarios: A, B, C, and D, with search durations of 6, 12, 18, and 24 hours respectively. We measured the Area Under the Detection Curve (AUDC) to assess the model's overall detection performance, and compared the evaluation results to the AUDC achieved during training. Additionally, we measured the time to first discovery, which indicates how quickly agents using the trained model can locate missing survivors during a disaster search and rescue. Lastly, we measured the total number of survivors found at the end of each mission.

1) AUDC per Scenario:

Figure 4(a) shows the AUDC values across the four evaluation scenarios. The results show how well the trained model generalizes under varying search durations. Scenario B (12h) achieved the highest AUDC of approximately 0.82, followed closely by Scenario A (6h) with a mean AUDC of 0.78. These results suggest that the model performs best under moderate search durations where agents have sufficient time to explore without excessive drift or redundancy. Scenario D (24h) recorded a mean AUDC of around 0.72, indicating a slight drop in detection consistency over extended missions. On the other hand, Scenario C (18h) had the lowest AUDC of approximately 0.65, with a wide confidence interval, reflecting high variability in agent performance across runs. This drop may be attributed to increased search space complexity and diminishing returns in survivor detection as time progresses.

When comparing the evaluation results to the overall training performance, the model shows consistent generalization

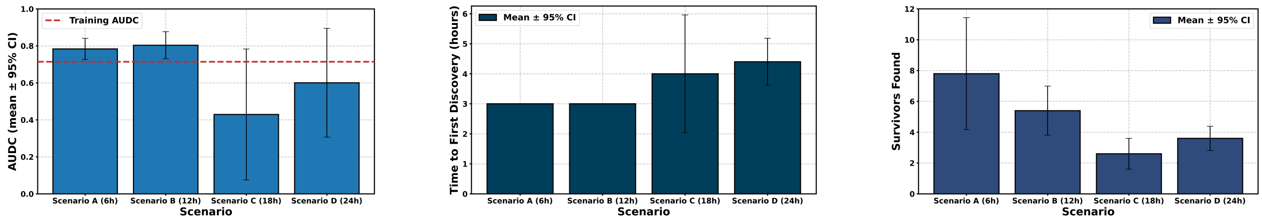

(a) AUDC per Scenario. (b) Time to First Discovery. (c) Overall Survivors Found.

Fig. 4. Model evaluation

as shown in Fig. 4(a). The results confirms that the training AUDC was approximately 0.75, which aligns closely with the evaluation values in Scenarios A and B. Scenario B exceeded the training performance, while Scenario A matched it. This comparison confirms that the model retains its detection ability when deployed, especially in scenarios that resemble the training conditions. The slight under performance in longer scenarios suggests that the model may benefit from additional training exposure to extended search durations or more diverse survivor distributions.

2) Time to First Discovery: Figure 4(b) highlights the time to first discovery metric which shows how quickly agents using the trained model can locate survivors. According to the results, in Scenario A, agents found the first survivor within approximately 3.0 hours, which is consistent with Scenario B, also averaging around 3.0 hours (reported times reflect simulation time, not real-world wall-clock time.) Scenario C recorded a higher mean of 4.0 hours, while Scenario D reached 4.5 hours. The narrow confidence intervals in Scenarios A and B indicate stable early detection, whereas the wider intervals in Scenarios C and D reflect inconsistent agent paths and delayed convergence. These results show that the model enables prompt detection in shorter missions but may require reinforcement strategies to enable more faster detection in extended operations.

3) Overall Survivors Found: Lastly, the number of survivors found at the end of each evaluation run provides a direct measure of search effectiveness. As shown in 4(c), Scenario A achieved the highest average with approximately 8 survivors found, followed by Scenario B with around 5. Furthermore, Scenario C recorded the lowest count at 3, and Scenario D slightly improved to 4. The decreasing trend from Scenario A to C reflects the challenge of maintaining search efficiency as mission time increases. The slight recovery in Scenario D may be due to agents eventually covering more ground, but the overall count remains below optimal. These results confirm that the trained model is most effective in shorter missions where survivor density and agent coordination are more favorable.

V. Conclusion

In this paper, we proposed a Federated DDQN for edge-based disaster response to improve search and rescue operations and ensure that missing survivors are located in a timely manner. Specifically, we adopt a Dueling Double Deep Q-Network (DDQN) that enables each first responder's device to act as an intelligent agent capable of learning optimal search patterns and strategies over time.

Our performance evaluation confirms that the trained model is most effective in shorter missions where survivor density and agent coordination are more favorable, and confirms strong detection capability in short and moderate missions.

In future work, we will extend the system through model fine-tuning, exposure to larger grid environments, and quantum aggregation to further shorten learning cycles and enhance robustness. Additionally, we will conduct extensive evaluation of X-FACE in disaster response using system-level simulations under diverse wireless network conditions.

Acknowledgment

This work was supported by Council for Science, Technology and Innovation(CSTI), Crossministerial Strategic Innovation Promotion Program (SIP), "Development of a Resilient Smart Network System against Natural Disasters" Grant Number JPJ012289 (Funding agency:NIED).

References

[1] T. Sakano, B. Ojetunde, A. Nakamura, and T. Fukada, "Proposal of a portable ICT system integrated with AI processing for enhancing disaster front-line operation," Disaster Resilience and Sustainable Development (DRSD 2023), Dec. 2023.

[2] Y. Yameng, Z. Lieping, S. Xiaoxu, W. Yilin, P. Jiansheng, amd Z. Jianchu, "Improved Double Deep Q Network Algorithm Based on Average Q-Value Estimation and Reward Redistribution for Robot Path Planning," Computers, Materials and Continua, Vol. 81, Issue 2, 2024, pp. 2769-2790.

[3] N. Merkle, R. Bahmanyar, C. Henry, S. M. Azimi, X. Yuan, S. Schopferer, V. Gstaiger, S. Auer, A. Schneibel, M. Wieland, and T. Kraf, "Drones4Good: Supporting Disaster Relief Through Remote Sensing and AI," https://doi.org/10.48550/arXiv.2308.05074.

[4] E. Samikwa, T. Voigt and J. Eriksson, "Flood Prediction Using IoT and Artificial Neural Networks with Edge Computing," 2020 International Conferences on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData) and IEEE Congr. on Cybermatics (Cybermatics), Rhodes, Greece, 2020, pp. 234-240.

[5] Stateczny, A.; Praveena, H.D.; Krishnappa, R.H.; Chythanya, K.R.; Babysarojam, B.B. "Optimized Deep Learning Model for Flood Detection Using Satellite Images," Remote Sens. 2023, 15, 5037.

[6] Alsumayt A, El-Haggar N, Amouri L, Alfawaer ZM, Aljameel SS. "Smart Flood Detection with AI and Blockchain Integration in Saudi Arabia Using Drones," Sensors. 2023; 23(11):5148.

[7] Abid, S.K.; Sulaiman, N.; Chan, S.W.; Nazir, U.; Abid, M.; Han, H.; Ariza-Montes, A.; Vega-Muñoz, A. "Toward an Integrated Disaster Management Approach: How Artificial Intelligence Can Boost Disaster Management," Sustainability 2021, 13, 12560.

[8] N. Andreassen, O. J. Borch, and A. K. Sydnes, "Information sharing and emergency response coordination," Safety Science, 130, 2020.