Urban Seismic Resilience: A Testbed for Studying Coordination and Communication in Multi-Agent Systems

Adyasha Patra **Satyam Srivastava** **Jay Chaudhary** **Abhay Jain**
Keshav Gupta **Seemandhar Jain**
University of California, San Diego
`{adpatra, sas059, jgchaudhary, arj012, keg019, sejain}@ucsd.edu`

Abstract

Large-scale disasters such as earthquakes create highly dynamic environments where rescue agents must operate under uncertainty, limited communication, and cascading hazards. We introduce an urban seismic disaster simulation platform designed for studying coordination strategies in multi-agent emergency response systems. The simulator models earthquakes, building collapses, fires, and aftershocks while supporting heterogeneous agents including scouts, medics, and firefighters. We evaluate a hierarchical architecture with a central command agent. Through controlled experiments across varying disaster densities and communication reliability levels, we observe that hierarchical coordination substantially improves rescue efficiency and survival outcomes. In low-density disaster scenarios, hierarchical agents achieve up to 7.6× higher survival rates compared to decentralized policies. We further analyze system robustness under message dropout and demonstrate graceful degradation, where local policies maintain operational performance even under severe communication failures. These results highlight the importance of hybrid coordination strategies combining centralized planning with local autonomy for resilient disaster response systems.

1 Introduction

Earthquakes produce complex urban disaster environments characterized by structural collapse, fires, infrastructure failure, and injured civilians requiring urgent rescue. Effective response requires coordination among heterogeneous responders operating in uncertain and rapidly evolving conditions. Traditional emergency response planning relies heavily on human coordination, but recent advances in autonomous agents and artificial intelligence present opportunities to augment disaster response with coordinated multi-agent systems.

However, evaluating coordination strategies for such systems remains difficult due to the lack of realistic and controllable simulation environments. Unlike what current environments can simulate, disaster environments involve cascading hazards, incomplete information, and communication constraints that challenge many existing multi-agent benchmarks.

To address this gap, we develop a configurable urban seismic disaster simulator that models earthquake events and their secondary effects, including structural collapses, fires, and aftershocks. Since the simulator is fully configurable and built from scratch, researchers can design their own disaster scenarios and experiments. This is a challenging research testbed for multi-agentic systems, which can test multiple abilities like planning, coordination, memory and reasoning.

Using this environment, we study a hierarchical agentic architecture in which a central commander agent, implemented as a large language model (LLM), performs high-level decision making and

Preprint.

task allocation for a team of field agents operating in the disaster environment. The field agents are responsible for navigating the environment and executing assigned tasks. While they primarily follow instructions issued by the commander, they also rely on lightweight heuristic policies to take autonomous actions when necessary.

We use this architecture to evaluate both the difficulty of the disaster scenarios generated by our simulator and the effectiveness of hierarchical coordination strategies in such environments. Our results provide preliminary insights into which coordination architectures may be better suited for disaster response scenarios characterized by uncertainty and rapidly evolving hazards.

Experimental results show that the hierarchical architecture improves survival rates by up to $7.6\times$ compared to decentralized field agents in low-density disaster scenarios. Additionally, we observe that disaster environments often favor fast local decision making over prolonged exploration, highlighting the importance of combining centralized planning with reactive local policies.

2 Related Work

Existing AI simulation environments have significantly advanced research in multi-agent reinforcement learning and embodied AI. However, they often lack the domain-specific dynamics required for studying complex urban disaster response scenarios. Widely used platforms such as OpenAI Gym [Brockman et al., 2016] and PettingZoo [Terry et al., 2021] provide flexible environments for reinforcement learning and multi-agent experimentation, but they do not natively model realistic urban disaster dynamics or coordinated multi-agent rescue tasks.

Several disaster-oriented simulation environments have also been developed, though they present different limitations. For instance, RESenv [Sun et al., 2023] focuses on highly realistic earthquake visualization through a game-engine-based environment, emphasizing graphical realism rather than supporting modern AI-driven experimentation. RoboCup Rescue [Kitano and Tadokoro, 1999] has long served as a benchmark for multi-agent disaster response research; however, it was originally designed for rule-based robotic agents and does not provide native integration with contemporary AI and machine learning workflows.

More recent research has explored the use of AI agents and large language models for disaster-resilient infrastructure and recovery planning (e.g., Park et al. [2023]). While these approaches demonstrate the potential of LLM-driven decision-making in complex environments, they often assume static damage conditions and perfectly reliable communication channels. Other simulation platforms, such as DisasterSim [USC Institute for Creative Technologies, 2013] and AI2-THOR [Kolve et al., 2017], either provide limited support for modeling cascading hazards or are not designed for large-scale multi-agent coordination in urban disaster environments.

To address these limitations, we introduce a configurable research testbed specifically designed to study how LLM-based agents coordinate, communicate, and make decisions under dynamic and high-stress disaster conditions.

3 Methodology

To study coordination strategies in disaster response scenarios, we design a configurable urban seismic disaster simulation environment and evaluate multi-agent coordination architectures within it. This section describes the environment design, disaster dynamics, agent roles, coordination framework, communication model, and evaluation methodology used in our experiments.

3.1 Urban Disaster Simulation Environment

The proposed simulator models earthquake-driven disasters occurring within a simplified urban environment represented as a two-dimensional grid. Each grid cell corresponds to a discrete location in the city and may contain buildings, roads, hazards, or civilians. The environment evolves dynamically over time as earthquake effects propagate and rescue agents interact with the environment.

City layouts are randomly initialized at the beginning of each simulation episode. Buildings are distributed across the grid with varying densities, which can configured, to simulate different urban conditions.

During each simulation step, the environment updates hazard states, agents execute actions, and the overall system state evolves accordingly. The simulator is designed to support experimentation with different disaster intensities and environmental conditions. Parameters such as earthquake magnitude, building density, fire propagation rates, and communication reliability can be adjusted to create a wide range of disaster scenarios.

3.2 Disaster Dynamics

At the beginning of every simulation, an earthquake is initiated with configurable magnitude, epicenter location, and a decay factor controlling the spatial attenuation of seismic intensity across the map. The intensity at any location is modeled using an exponential decay function:

$$I = I_0 e^{-kr} \tag{1}$$

where I denotes the earthquake intensity at a given location, I_0 represents the initial intensity at the epicenter, r is the distance from the epicenter, and k is a decay constant that determines how rapidly the seismic intensity diminishes with distance.

In subsequent simulation steps, the earthquake event triggers several cascading hazards that represent common consequences observed in real-world urban earthquakes, including structural collapses, fires, and road blockages that affect emergency response coordination.

Building Collapses: At the onset of an earthquake event, a subset of buildings collapse based on the earthquake intensity and structural vulnerability parameters. Every building has an integrity paramter which declines according to the damage caused by the earthquake on the building cells, and when the integrity reaches below a collapse threshold, the building probabilistically collapses on the original building cells as well as the neighboring cells in the grid. Collapsed buildings may trap civilians who require rescue operations from field agents.

Fire Outbreaks: Fires may ignite in damaged buildings due to gas leaks, electrical faults, or structural failures. Once initiated, fires can spread to neighboring structures depending on configurable fire propagation probabilities.

Aftershocks: The simulator also models aftershock events that can be tuned with the required time step of initiaition and magnitude. Aftershocks may trigger additional building collapses and exacerbate ongoing hazards, forcing rescue agents to adapt to continuously changing environmental conditions.

Any civilian trapped in a collapsed building or in a building with fire is termed as a *victim*. Victims have a decay factor on their lives, which is configurable. They have to be rescued in their survival time window.

3.3 Agent Roles and Capabilities

The rescue system consists of multiple heterogeneous agents, each representing a specialized emergency response unit. Field agents operate within the environment using a combination of local observations, and communication and task assignments from the commander agent.

Commander Agent

We model the commander agent using an LLM. At every time step, it has three things in its context.

- Mental Map: Initialized with the ground truth map, before the earthquake hit, this shows the city layout to the commander. As field agents explore and discover hazards(fire and collapsed buildings) and victims, the commander gets updated and the respective cells in the mental map get updated.
- Zone summaries: We give the commander access to surveillance data. We model it by creating zones, which are 10x10 subgrids within the ground truth map. For every zone, we give the commander estimates of the number of collapsed buildings, number of fires and number of trapped victims, without revealing their precise locations.
- Field agents' locations and availability: We assume that the field agents are GPS tracked. So the commander always has track of their precise locations. Additionally, when a field agent

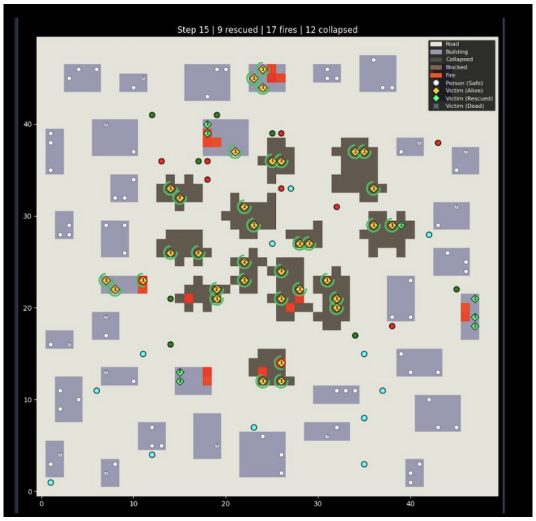

Figure 1: Snapshot of the urban disaster simulation environment following an earthquake event, illustrating damaged structures, fires, and civilians requiring rescue.

completes an allocated task, it sends a report back to the commander using which free and busy agents are tracked. The commander gets a list of all available agents at every time step.

Using all of the above context and details about the powers of each type of field agent, the commander allocates tasks to the available agents. The task is allocated using the exact location/zone that the specific agent should go to, including the path calculated using A* algorithm but only with access to the mental map of the commander, and not the actual ground truth map.

Three primary field agent roles implemented are:

Scout Agents: Scout agents are responsible for exploring the environment and identifying hazards or victims in unknown regions of the environment. They maintain a local observation radius that allows them to detect nearby hazards, collapsed buildings, and victims, and their precise locations.

Medic Agents: Medic agents perform rescue operations by stabilizing injured civilians in buildings which have collapsed.

Firefighter Agents: Firefighter agents are responsible for mitigating environmental hazards by extinguishing fires. By controlling fires, firefighters prevent further structural damage and reduce the likelihood of additional civilian injuries.

We model our simulation such that while completing an assigned task, the medic and firefighter agents can save any victims in their immediate redius on the way to their destination.

Each agent type has distinct capabilities and action spaces, requiring effective coordination to ensure that resources are allocated appropriately across the environment. The number of each type of field agent is also configurable.

3.4 Communication Model

Communication between agents is modeled as a message-bus system with configurable reliability. Agents can transmit observations, hazard reports, and status updates to the central command unit. To simulate realistic disaster conditions, the simulator introduces message dropout probabilities that represent communication failures caused by damaged infrastructure or network congestion. Field agents therefore maintain local fallback policies that allow them to continue operating autonomously when command instructions are unavailable.

4 Experiments and Results

We evaluate the proposed multi-agent coordination framework using the urban seismic disaster simulator described in Section 3. Each configuration is executed across randomized disaster scenarios, and results are reported as the mean and variance across 10 independent simulations.

4.1 Experimental Setup

All experiments are performed in a grid-based urban disaster environment with procedurally generated city layouts and configurable disaster parameters. Disaster severity is controlled by varying building density, earthquake intensity, and civilian survival windows. Higher building density increases the probability of structural collapse, while stronger earthquake intensities generate more widespread cascading hazards across the environment. These variations allow evaluation under low-, medium-, and high-disaster conditions. The commander agent in the hierarchical architecture is implemented using the *gpt-oss-120B* language model, which performs high-level reasoning, planning, and task allocation based on observations reported by field agents.

4.2 Evaluation Metrics

We use two primary metrics.

Survival Rate. Fraction of civilians rescued within the simulation time horizon.

Mental Map Fidelity. Similarity between the commander's internal map and the ground-truth environment, measured as the fraction of correctly predicted grid cells.

4.3 Experiment 1: Hierarchical vs Decentralized Coordination

RQ1: How does hierarchical coordination compare to decentralized decision making?

Setup We compare two coordination strategies. In the *hierarchical architecture*, a central commander aggregates scout observations, maintains a global mental map, and assigns tasks to field agents. In the *decentralized architecture*, agents operate independently using only local observations within a limited sensing radius.

Results Hierarchical coordination significantly outperforms decentralized coordination across disaster settings. In low-density environments, hierarchical coordination achieves approximately **84% survival rate**, compared to **11%** for decentralized agents (a **7.6× improvement**).

The advantage persists across increasing disaster densities: survival rates decrease to **51% vs 18%** in medium-density scenarios and **30% vs 18%** in high-density environments. The performance gap narrows at high density because environmental constraints such as travel time and simultaneous hazards limit the impact of coordination.

Additionally, hierarchical coordination improves global situational awareness, increasing mental map fidelity from approximately **48% in decentralized agents to about 62%** under centralized coordination.

4.4 Experiment 2: Communication Dropout Resilience

RQ2: How resilient is the coordination system under communication failures?

Setup We simulate degraded communication infrastructure by introducing probabilistic message dropouts between the commander and field agents.

Results As message dropout probability increases, performance gradually degrades due to reduced accuracy in the commander's global state estimation. Missing scout reports lead to inconsistencies in the global mental map, causing redundant exploration and delayed rescue actions.

Despite these effects, the system demonstrates **graceful degradation**. Even under severe communication loss, agents maintain a survival rate floor of approximately **54%** due to their local reactive behaviors. Mental map fidelity decreases moderately under high dropout conditions but remains above **50%**, indicating that the system can still maintain a usable global representation.

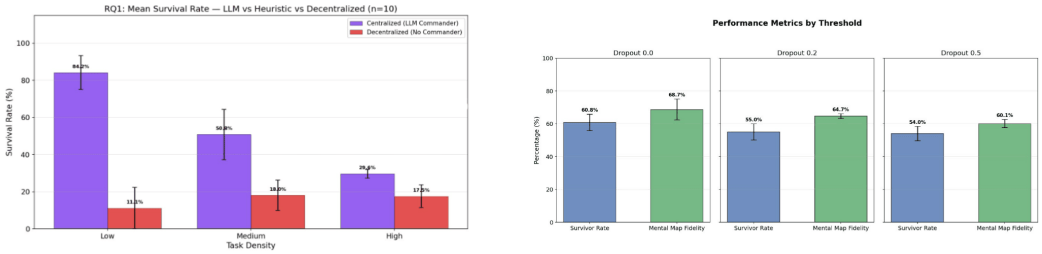

(a) Hierarchical vs Decentralized

(b) Communication Dropout

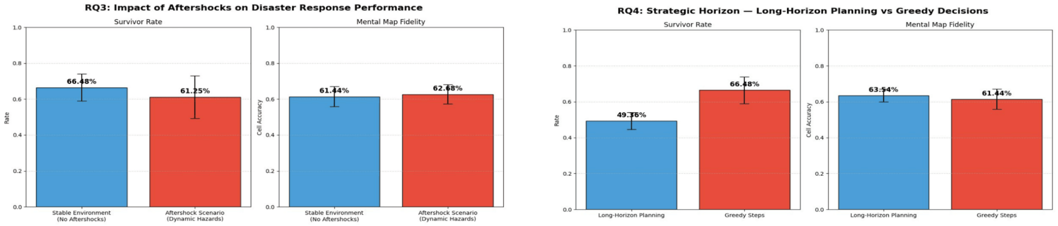

(c) Dynamic Disaster Adaptation

(d) Greedy vs Long-Horizon Planning

Figure 2: Experimental results across four coordination experiments. (a) Hierarchical coordination significantly improves survival rates compared to decentralized agents. (b) System performance degrades gradually under communication dropouts but maintains baseline operation. (c) Agents adapt to aftershock events with only moderate performance reduction. (d) Greedy decision policies outperform long-horizon planning in dynamic disaster environments.

4.5 Experiment 3: Adaptation to Dynamic Disaster Events

RQ3: Can the system adapt to disruptive events during ongoing rescue operations?

Setup To evaluate adaptability, we introduce stochastic aftershock events during rescue operations. Aftershocks trigger additional building collapses and fires, forcing agents to update situational awareness and adjust task priorities.

Results Aftershock events produce a moderate reduction in rescue performance. Civilian survival decreases from approximately **66% to 61.4%**, corresponding to a **4.6% drop**. Despite these disruptions, the system remains stable.

Importantly, the commander's mental map fidelity remains relatively stable at approximately **62%**, suggesting that the coordination system effectively incorporates new observations. The performance reduction therefore appears to be driven primarily by increased environmental hazards rather than failures in situational awareness.

4.6 Experiment 4: Greedy vs Long-Horizon Planning

RQ4: Which strategy performs better: greedy reactive dispatch or long-horizon planning?

Setup We compare two commander policies. The *greedy policy* prioritizes immediate high-value actions such as rescuing nearby civilians or extinguishing active fires. The *long-horizon policy* attempts to optimize global objectives such as exploration coverage by scout agents for more steps before assigning tasks to medic and firefighter agents.

Results Greedy strategies outperform long-horizon planning in dynamic disaster environments. Greedy policies achieve approximately **66% survival rate**, while long-horizon planning achieves about **49%**. This corresponds to a **34.7% relative improvement** in rescue effectiveness.

Interestingly, long-horizon planning slightly improves environmental observability, achieving **63.5% map fidelity** compared to **61.4%** under greedy planning. However, the additional planning overhead delays rescue actions, reducing overall survival outcomes.

5 Conclusion

We presented an urban seismic disaster simulation platform for studying coordination strategies in multi-agent rescue systems. Using this environment, we evaluated a hierarchical coordination architecture in which a central commander assigns tasks to specialized field agents.

Our experiments show that hierarchical coordination significantly improves rescue efficiency by enabling global situational awareness and structured task allocation, achieving up to a 7.6× improvement over decentralized setups in civilian survival rates in certain scenarios. We also observe graceful degradation under communication failures, where autonomous fallback behaviors allow agents to maintain operational effectiveness despite message dropouts. Finally, our results indicate that reactive greedy decision strategies outperform long-horizon planning approaches in dynamic disaster environments where rapid response is critical.

These findings highlight the importance of hybrid coordination architectures that combine centralized planning with decentralized autonomy. The simulator introduced in this work provides a flexible testbed for future research on agentic AI systems and resilient multi-agent coordination in complex disaster response scenarios.

6 Limitations

While the proposed simulation framework and hierarchical coordination architecture provide useful insights into multi-agent disaster response, several limitations should be acknowledged.

Although each experiment was repeated ten times and results are reported as mean ± variance, this number of runs provides only a moderate level of statistical confidence. Increasing the number of experimental runs and evaluating across a broader range of disaster configurations would further strengthen the reliability of the reported trends.

Our experiments use a single large language model to implement the commander agent. While we expect similar qualitative trends across different models, it remains uncertain how model-specific inductive biases may influence planning behavior, coordination strategies, and decision-making performance in hierarchical multi-agent settings.

In the current implementation, field agents rely on heuristic policies within the hierarchical framework. This design allows us to isolate the effects of centralized coordination, but it limits the autonomy and reasoning capabilities of individual agents. Replacing these heuristic policies with fully autonomous LLM-based agents may improve adaptability and robustness, particularly in scenarios involving communication failures or message dropout.

Finally, the simulated urban environments use randomly generated city layouts rather than real-world geographic maps. While this enables flexible experimentation, synthetic layouts may not fully capture the structural constraints and infrastructure patterns present in real cities. Extending the simulator to incorporate realistic urban layouts would require additional development effort and may influence agent interactions, navigation behavior, and overall system dynamics.

References

Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym, 2016.

Hiroaki Kitano and Satoshi Tadokoro. Robocup rescue: Search and rescue for large-scale disasters. In *IEEE International Conference on Systems, Man, and Cybernetics*, 1999.

Eric Kolve, Roozbeh Mottaghi, Winson Han, Luke VanderBilt, Luca Weihs, Alvaro Herrasti, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. Ai2-thor: An interactive 3d environment for visual ai. In *arXiv preprint arXiv:1712.05474*, 2017.

Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. *Proceedings of the ACM CHI Conference*, 2023.

Yitong Sun, Hanchun Wang, Zhejun Zhang, Cyriel Diels, and Ali Asadipour. Resenv: A realistic earthquake simulation environment based on unreal engine, 2023. URL `https://arxiv.org/abs/2311.07239`.

Justin K. Terry, Benjamin Black, Nathaniel Grammel, Mahanth Jayakumar, Ananth Hari, and et al. Pettingzoo: Gym for multi-agent reinforcement learning. *NeurIPS Datasets and Benchmarks*, 2021.

USC Institute for Creative Technologies. Disastersim, 2013. URL `https://ict.usc.edu/research/projects/disastersim/`. Game-based disaster response training simulation.