

1 Adaptive Disaster Relief Routing Using Deep Reinforcement Learning Under 2 Dynamic Road Disruptions

3 Natural disasters such as floods frequently disrupt transportation networks, significantly hindering
4 emergency relief logistics and delaying the delivery of critical supplies. Traditional routing optimization
5 methods, including shortest-path algorithms and metaheuristic approaches, typically assume static
6 network conditions and struggle to adapt to rapidly evolving road disruptions. This paper proposes a Deep
7 Reinforcement Learning (DRL) framework for adaptive disaster relief routing in dynamically changing
8 environments. The problem is formulated as a Dynamic Vehicle Routing Problem (DVRP) modeled using
9 a Markov Decision Process (MDP). A grid-based city simulation with nine demand nodes and stochastic
10 road blockages is developed to emulate flood-induced transportation disruptions. Two DRL algorithms,
11 Deep Q-Network (DQN) and Proximal Policy Optimization (PPO), are evaluated and compared against a
12 Genetic Algorithm (GA) baseline. Experimental results demonstrate that the PPO-based routing policy
13 achieves improved adaptability and reduced delivery delays under high-disruption scenarios,
14 outperforming traditional optimization methods. Furthermore, the DRL agents exhibit faster recovery and
15 policy adjustment following sudden road blockages. The findings highlight the potential of reinforcement
16 learning-based decision systems for real-time disaster logistics planning and adaptive emergency
17 response.

18 **Keywords:** Disaster Management, Deep Reinforcement Learning, Dynamic Vehicle Routing,
19 Emergency Logistics, Proximal Policy Optimization (PPO), Deep Q-Network (DQN), Disaster Response
20 Systems.

21
22

23 Introduction

24 Natural disasters such as floods, earthquakes, and hurricanes frequently disrupt transportation
25 infrastructure and severely affect emergency relief operations. Rapid and efficient delivery of essential
26 supplies, including food, medical aid, and rescue equipment, is critical for minimizing loss of life and
27 ensuring effective disaster response. However, transportation networks in disaster-affected regions often
28 experience dynamic disruptions such as blocked roads, damaged bridges, and unpredictable traffic
29 conditions.

30 Traditional routing methods such as Dijkstra's shortest path algorithm, A*, and metaheuristic techniques
31 including Genetic Algorithms (GA) and Ant Colony Optimization assume relatively static environments.
32 These methods require repeated recomputation when network conditions change, which limits their
33 effectiveness in dynamic disaster scenarios.

34 Recent advancements in Artificial Intelligence, particularly Deep Reinforcement Learning (DRL), have
35 demonstrated strong capabilities in sequential decision-making problems under uncertainty. Unlike static
36 optimization approaches, DRL agents can learn adaptive routing policies that respond to environmental
37 changes without requiring complete re-optimization.

38 In this research, disaster relief routing is modeled as a Dynamic Vehicle Routing Problem (DVRP) under
39 stochastic road disruptions. A grid-based simulation environment representing flood-affected urban areas
40 is developed. Two DRL algorithms, Deep Q-Network (DQN) and Proximal Policy Optimization (PPO),
41 are evaluated and compared with a Genetic Algorithm baseline.

42
43
44
45
46
47

Table 1: Performance Comparison of Machine Learning Algorithms in Disaster Management

48 **Deep Network Architecture**

49 **1. Input Modalities**

Algorithm	Study / Author	Accuracy (%)	RMSE	Remarks
Linear Regression	Yasmin et al. (2025)	78.4	0.42	Baseline, weak for nonlinear disaster data
Random Forest	Kim & Choi (2025)	89.7	0.28	Good for multispectral & satellite data
XGBoost	Raj et al. (2025)	92.3	0.22	Strong for multimodal and tabular data
SVM	Singh et al. (2022)	84.1	0.33	Good for risk classification
LSTM	Anusha (2023)	94.8	0.18	Best for time-series forecasting
CNN	Zhang et al. (2022)	91.2	0.25	Highly effective for damage detection
CNN-LSTM	Martinez et al. (2023)	96.1	0.15	Best overall for spatiotemporal modeling

50 The model processes three major data sources:

51 i. Satellite / UAV images → spatial features

52 ii. Weather & environmental time-series → temporal patterns

53 iii. Population, infrastructure & historical disaster data → static features

54 To handle these, ResQ uses a hybrid deep learning network combining CNN, LSTM, and a Fusion Dense
55 Network.

56

57 **2. Architecture Overview**

58 **(A) CNN Module – Spatial Feature Extractor**

59 Used for damage assessment, flood extent detection, cyclone impact zones.

60 Layers:

61 A. Conv2D (32 filters, 3×3)

62 B. ReLU + Batch Normalization

63 C. MaxPooling

64 D. Conv2D (64 filters, 3×3)

65 E. ReLU + Dropout

66 F. Flatten → Spatial Feature Vector

67 This follows architectures used by Kim & Choi (2025) and Zhang et al. (2022).

68 **(B) LSTM Module – Temporal Sequence Model**

69 Used for forecasting demand, rainfall patterns, disaster progression.

70 Input: Weather parameters across time (rainfall, humidity, wind speed, water level)

71 Layers:

72 I. LSTM (128 units)

73 II. Dropout

74 III. LSTM (64 units)

75 IV. Dense layer → Temporal Feature Vector

76 Similar to Anusha (2023) and ISRO flood forecasting models.

77 **(C) Static Feature Processing Module**

78 For population density, road accessibility, historical damage patterns.

79 Layers:

80 a) Dense (64)


81 b) ReLU

82 c) Dense (32)

83 ReLU → Static Feature Vector

84 **(D) Multimodal Fusion Layer**

85 All 3 feature vectors are concatenated:



Architecture of Our CNN System

86

87 Fusion block layers:

88

89 1) Dense (128), ReLU

90 2) Batch Normalization

91 3) Dense (64), ReLU

92 4) Dropout

93 5) Dense (1) → Final Output


94 6) demand prediction / damage level / resource allocation score

95

96 This approach follows multimodal fusion strategies from Martinez et al. (2023) and Raj et al. (2025).

97

An Example Demonstrating the Idea of
Multilevel CNN Networks



98

99 Methodology

100 A. Simulation Environment

101 A grid-based city environment is constructed to simulate disaster scenarios. The grid size is set to 15×15
102 nodes, representing urban road intersections.

103 The environment includes:

- 104 ● One depot location
- 105 ● Nine demand nodes
- 106 ● Dynamic road blockages

107 Flood disruptions randomly block 10–30% of roads at each episode. Additionally, a sudden disruption
108 event occurs mid-episode to simulate unexpected infrastructure damage.

109

110 B. Genetic Algorithm Baseline

111 The Genetic Algorithm (GA) is used as a traditional optimization baseline.

112 Key components include:

- 113 ● Chromosome representation as routing sequences
- 114 ● Fitness function based on total travel distance
- 115 ● Selection, crossover, and mutation operators

116 GA computes an optimized route but requires recomputation when road disruptions occur.

117

118 C. Deep Q-Network

119 Deep Q-Network (DQN) is used as a value-based reinforcement learning algorithm. The network
120 approximates the action-value function:

121

$$Q(s, a)$$

122

123

124 The loss function for training is:

125

$$\mathbf{L}(\theta)=\mathbf{E}[(\mathbf{r}+\gamma\max_{\mathbf{a}'}\mathbf{Q}(\mathbf{s}',\mathbf{a}')-\mathbf{Q}(\mathbf{s},\mathbf{a}))^2]$$

126

Experience replay and target networks are used to stabilize learning.

127

128 **D. Proximal Policy Optimization**

129

Proximal Policy Optimization (PPO) is a policy-gradient method that uses an actor-critic architecture.

130

PPO improves training stability through clipped objective functions:

131

$$\mathbf{LCLIP}(\theta)=\mathbf{E}[\mathbf{min}(\mathbf{rt}(\theta)\mathbf{A}t,\mathbf{clip}(\mathbf{rt}(\theta),\mathbf{1}-\epsilon,\mathbf{1}+\epsilon)\mathbf{A}t)].$$

132

PPO enables robust learning in dynamic environments.

133

134 **E. Experimental Setup**

135

The experiments are conducted using a custom simulation environment implemented in Python.

136

Key experimental settings include:

137

Parameter	Value
Grid Size	15×15
Demand Nodes	9
Training Episodes	15,000
Disruption Rate	10–30%
Algorithms	GA, DQN, PPO

144

The experiments were implemented using:

- 145 ● Python 3.10
- 146 ● PyTorch
- 147 ● Stable-Baselines3
- 148 ● OpenAI Gym environment

149 Performance metrics include:

- 150 ● Average delivery delay
- 151 ● Success rate
- 152 ● Adaptation time after disruption
- 153 ● Total travel distance

154

155 **Discussion**

156 The experimental evaluation highlights the effectiveness of reinforcement learning-based routing
157 strategies in dynamically disrupted transportation environments. The comparative analysis between the
158 Genetic Algorithm (GA), Deep Q-Network (DQN), and Proximal Policy Optimization (PPO)
159 demonstrates that learning-based routing policies provide significant advantages when operating under
160 uncertain disaster conditions.

161 One of the primary observations from the experiments is the improved adaptability of reinforcement
162 learning agents compared to traditional optimization approaches. The Genetic Algorithm baseline
163 performs well in static routing scenarios but struggles in environments where road conditions change
164 frequently. This limitation arises because the algorithm must recompute optimal routes whenever
165 disruptions occur, leading to increased computational overhead and delayed decision-making. In contrast,
166 reinforcement learning agents learn adaptive policies during training, enabling them to respond to
167 environmental changes without requiring complete re-optimization.

168 Among the reinforcement learning algorithms evaluated, PPO exhibited the most stable and efficient
169 performance. The policy-gradient framework used in PPO allows the agent to directly learn a stochastic
170 policy that balances exploration and exploitation. The clipped objective function used during training
171 helps prevent excessive policy updates, resulting in more stable learning behavior. This stability becomes
172 particularly important in disaster environments where sudden disruptions can cause large changes in the
173 state space.

174 The Deep Q-Network algorithm also demonstrated strong performance compared to the Genetic
175 Algorithm baseline. However, the convergence rate of DQN was slower than that of PPO, primarily due
176 to the value-based learning framework and the challenges associated with estimating optimal Q-values in
177 large state spaces. In highly dynamic environments, DQN occasionally required additional training
178 episodes to achieve stable routing policies. Despite these limitations, DQN still provided significant
179 improvements in adaptability compared to traditional optimization methods.

180 Another important observation is the impact of disruption intensity on routing performance. As the
181 percentage of blocked roads increased from 10% to 30%, the overall difficulty of the routing problem
182 increased significantly. Traditional optimization methods experienced larger performance degradation

183 under higher disruption levels. In contrast, reinforcement learning agents demonstrated more robust
184 behavior, maintaining relatively stable performance even under severe disruptions. This result suggests
185 that DRL-based routing policies are better suited for environments where infrastructure damage evolves
186 over time.

187 The mid-episode disruption event introduced in the simulation environment further illustrates the
188 advantages of reinforcement learning methods. When sudden road blockages occurred during the routing
189 process, the reinforcement learning agents were able to adapt their routing strategies quickly and recover
190 from performance drops within a few time steps. The PPO agent in particular demonstrated faster
191 recovery compared to DQN and GA. This behavior indicates that the learned policy effectively captured
192 the underlying structure of the environment and could generalize to previously unseen disruption patterns.

193 In addition to adaptability, reinforcement learning methods also provide benefits in terms of scalability
194 and generalization. Once trained, the RL agent can apply the learned policy across multiple disaster
195 scenarios without requiring significant recomputation. This property is particularly useful in real-world
196 disaster response systems where time-critical decisions must be made under uncertainty. The ability to
197 rapidly generate routing decisions without extensive recomputation makes reinforcement learning
198 approaches highly suitable for emergency logistics planning.

199 Despite these advantages, several challenges remain when applying reinforcement learning to disaster
200 routing problems. First, the performance of DRL models is highly dependent on the quality of the training
201 environment. If the simulation environment does not accurately represent real-world disaster conditions,
202 the learned policies may not generalize effectively. Second, reinforcement learning methods require
203 significant training time and computational resources, especially when dealing with large-scale routing
204 problems. Finally, explainability remains an important concern when deploying AI-based decision
205 systems in critical applications such as disaster management.

206 Future research can address these challenges by incorporating more realistic disaster simulations,
207 including real-world geographic data and traffic information. The integration of graph neural networks
208 may further improve the ability of reinforcement learning models to capture complex network structures
209 in large urban environments. Additionally, multi-agent reinforcement learning frameworks could enable
210 coordination among multiple emergency vehicles, improving overall disaster response efficiency.

211 Overall, the findings of this study demonstrate that reinforcement learning offers a promising direction for
212 adaptive disaster logistics systems. By enabling intelligent routing decisions in dynamic environments,
213 DRL-based approaches have the potential to significantly enhance emergency response operations and
214 improve the efficiency of disaster relief distribution.

215

216 **Conclusion**

217 This study presented a Deep Reinforcement Learning framework for adaptive disaster relief routing under
218 dynamic road disruptions. The problem was modeled as a Dynamic Vehicle Routing Problem and
219 evaluated using a grid-based simulation environment with stochastic road blockages.

220 Experimental results demonstrated that reinforcement learning methods outperform traditional
221 optimization techniques in dynamic disaster scenarios. In particular, the PPO-based routing policy

222 achieved lower delivery delays and faster adaptation to sudden infrastructure disruptions compared to the
223 Genetic Algorithm baseline. These findings highlight the potential of reinforcement learning for
224 improving real-time disaster logistics planning and emergency response operations.

225 Future research will focus on extending the proposed framework to multi-agent routing systems,
226 integrating real-world geographic datasets, and applying graph neural networks for large-scale disaster
227 management scenarios.

228 .

229 **References**

230 [1] **S. Kadyrov, A. Azamov, Y. Abdumajitov, and C. Turan**, “Deep reinforcement learning for dynamic
231 vehicle routing with demand and traffic uncertainty,” *Operations Research Perspectives*, 2025.

232 [2] **J. Hu, Y. Wang, and Z. Chen**, “A deep reinforcement learning–based route planning
233 framework for vehicle routing problems,” *Mathematics*, vol. 13, no. 13, 2025.

234 [3] **F. Pu, X. Zhao, and Y. Zhang**, “Recent advances in disaster emergency response planning,”
235 *Smart Energy Systems Review*, 2025.

236 [4] **K. Mili, A. Benkhelifa, and S. Al-Maadeed**, “Adaptive vehicle routing for humanitarian aid
237 in conflict and disaster environments,” *Frontiers in Future Transportation*, 2025.

238 [5] **H. Gong, L. Chen, and W. Li**, “Deep reinforcement learning for drone route optimization in
239 post-disaster road assessment,” *arXiv preprint*, 2025.

240 [6] **A. Konovalenko**, “Optimizing a dynamic vehicle routing problem with deep reinforcement
241 learning,” *Logistics*, vol. 8, no. 4, 2024.

242 [7] **Z. Iklassov, K. Arapkhanov, and S. Kim**, “Reinforcement learning for solving stochastic
243 vehicle routing problems,” *Proceedings of Machine Learning Research*, 2024.

244 [8] **W. Kool, H. van Hoof, and M. Welling**, “Deep policy dynamic programming for vehicle
245 routing problems,” *International Conference on Artificial Intelligence*, 2024.

246 [9] **P. Upadhyay, S. Roy, and A. Ghosh**, “An improved deep reinforcement learning routing
247 technique for dynamic networks,” *Scientific Reports*, vol. 13, 2023.

248 [10] **J. Zhang and T. Van Woensel**, “Dynamic vehicle routing with random requests: A
249 literature review,” *International Journal of Production Economics*, vol. 250, 2023.

250 [11] **M. Cappart, D. Chételat, E. B. Khalil, and A. Lodi**, “Combinatorial optimization and
251 reasoning with graph neural networks,” *Journal of Machine Learning Research*, vol. 24, 2023.

252 [12] **H. Su, Y. Liu, and Q. Wang**, “**EMVLight**: A multi-agent reinforcement learning
253 framework for emergency vehicle routing and traffic signal control,” *IEEE Conference on
254 Cyber-Physical Systems*, 2022.

255 [13] **E. Rodríguez-Esparza, J. J. Ríos-Mercado, and M. Laguna**, “A hyper-heuristic based on
256 reinforcement learning for the capacitated electric vehicle routing problem,” IEEE World
257 Congress on Computational Intelligence, 2022.

258 [14] **J. Li, Y. Li, and X. Liu**, “Deep reinforcement learning for solving heterogeneous
259 capacitated vehicle routing problems,” IEEE Transactions on Intelligent Transportation Systems,
260 2021.

261 [15] **A. Delarue, T. Dietterich, and J. Gonzalez**, “Reinforcement learning with combinatorial
262 actions,” Advances in Neural Information Processing Systems (NeurIPS), 2020.

263

264

UNDER PEER REVIEW IN IJAR