

A Multi-Modal Decision-Support Framework for Real-Time Urban Infrastructure Vulnerability Assessment and Emergency Resource Allocation During Extreme Weather Events

Authors

Judy Anna, Taloza Capul, Abill Robert, Ada John, Reslie Pagador, Sadi Iddi, Jaryi Sunio

Date; October 6, 2026

Abstract

Extreme weather events increasingly threaten urban infrastructure systems, yet emergency managers lack integrated tools that can simultaneously assess vulnerability and optimize resource allocation in real time. Existing approaches remain fragmented: hazard prediction models operate independently of resource optimization algorithms, and most systems rely on unimodal data streams that fail to capture the multidimensional nature of urban crises. This study develops and validates a multi-modal decision-support framework that integrates meteorological data, infrastructure sensor networks, and geospatial indicators to predict vulnerability and allocate emergency resources dynamically during flood events. The framework employs a hybrid deep learning architecture combining LSTM networks for temporal forecasting with gradient boosting for vulnerability classification, coupled with a stochastic optimization module for resource allocation. Validation using historical flood data from three urban regions demonstrates that the framework achieves 89.4% accuracy in identifying high-vulnerability infrastructure nodes within a six-hour lead time. The integrated approach reduces resource allocation latency by 42%

compared to static allocation methods and improves coverage equity by 28%. These findings establish that multi-modal integration significantly enhances both predictive accuracy and operational efficiency in emergency management. The framework provides a replicable model for municipalities seeking to strengthen urban resilience through data-driven decision support, with implications extending to other climate-sensitive infrastructure systems and emergency contexts.

Keywords: multi-modal decision support, urban vulnerability assessment, emergency resource allocation, extreme weather adaptation

1. Introduction

1.1 Background

Urban infrastructure systems—encompassing transportation networks, power distribution, water management, and communication systems—face unprecedented stress from increasingly frequent and intense extreme weather events. Climate change has amplified both the frequency and severity of precipitation-driven flooding, creating cascading failures that overwhelm traditional emergency response mechanisms [2]. Between 2015 and 2024, urban flooding alone caused an estimated \$47 billion in infrastructure damage across major metropolitan areas worldwide, with disproportionate impacts concentrated in low-lying, densely populated districts characterized by aging infrastructure and limited adaptive capacity [15].

The challenge confronting emergency managers is fundamentally dual: they must rapidly assess which infrastructure components are most vulnerable to imminent hazard, while simultaneously determining how to allocate scarce resources—personnel, equipment, and supplies—to mitigate damage and protect affected populations. These tasks are temporally coupled and computationally demanding, requiring synthesis of heterogeneous data streams under severe time constraints. Traditional approaches treat vulnerability assessment and resource allocation as sequential, independent activities, often producing suboptimal outcomes when hazard conditions evolve faster than assessment cycles [7].

Recent advances in artificial intelligence have created new possibilities for integrated emergency decision support. Multi-modal learning approaches that combine satellite imagery, ground-based sensors, and textual data streams have demonstrated superior predictive performance compared to unimodal methods [3][13]. Concurrently, optimization techniques including stochastic programming and multi-criteria decision analysis have improved resource allocation efficiency in humanitarian logistics [4][7]. However, these developments have proceeded largely in parallel, with limited integration between predictive analytics and prescriptive optimization within operational decision-support systems [1].

1.2 Problem Statement

Despite substantial investment in disaster prediction and emergency management technologies, a critical gap persists between predictive capability and operational decision-making. Current systems typically excel at either hazard forecasting or resource optimization, but few provide integrated decision support that translates vulnerability predictions into actionable allocation recommendations in real time [16]. This fragmentation creates several specific problems.

First, existing vulnerability assessment models often rely on static exposure indicators—such as historical flood maps or census-based demographic data—that fail to capture the dynamic evolution of hazard conditions during an event. While recent deep learning approaches have improved flood prediction accuracy [5][10], these models rarely incorporate infrastructure-specific vulnerability parameters or produce outputs directly usable by resource allocation algorithms.

Second, resource allocation methods in emergency management predominantly employ pre-disaster static allocation schemes that assume worst-case demand scenarios [7]. These approaches waste resources when actual conditions prove less severe than anticipated, while failing to adapt when conditions worsen unexpectedly. Dynamic allocation methods exist in theory but require real-time vulnerability assessments that current prediction systems cannot reliably provide at operational timescales.

Third, the integration challenge itself remains unresolved. Multi-modal data fusion techniques have been developed independently for remote sensing [13], sensor networks [3], and text analytics [18], but no validated framework exists that harmonizes these heterogeneous inputs into a unified decision-support architecture capable of serving emergency managers under operational constraints.

The unsolved issue is therefore clear: no validated, replicable framework exists that integrates multi-modal vulnerability prediction with dynamic resource allocation optimization in a form suitable for real-time emergency decision support during extreme weather events.

1.3 Objectives of the Study

General objective: To develop and validate a multi-modal decision-support framework that enables real-time urban infrastructure vulnerability assessment and dynamic emergency resource allocation during extreme weather events.

Specific objectives:

1. To identify and integrate key multi-modal predictors—including meteorological forecast data, infrastructure sensor readings, and geospatial vulnerability indicators—that most accurately predict infrastructure disruption during flood events.
2. To design a hybrid predictive architecture that combines temporal sequence modeling with spatial vulnerability classification to generate actionable vulnerability assessments at operational lead times.

3. To develop a dynamic resource allocation optimization module that translates vulnerability predictions into prioritized resource deployment recommendations under uncertainty.
4. To validate the integrated framework against historical flood events and compare its performance to conventional static allocation and unimodal prediction approaches.

1.4 Research Questions

1. What combination of multi-modal data sources and modeling techniques most accurately predicts urban infrastructure vulnerability during extreme weather events?
2. How does the proposed integrated framework compare to traditional methods in terms of prediction accuracy, resource allocation efficiency, and decision lead time?
3. What implementation barriers—technical, organizational, and data-related—must be addressed for operational deployment of multi-modal decision-support systems in municipal emergency management?

1.5 Significance of the Study

This research matters for several constituencies. For **emergency management practitioners and administrators**, the framework provides a validated tool that can reduce response latency and improve resource utilization during crises, potentially saving lives and reducing economic losses. The 89.4% vulnerability prediction accuracy and 42% reduction in allocation latency represent operationally meaningful improvements over current methods.

For **policymakers**, the study offers evidence-based guidance on investing in integrated decision-support infrastructure. The finding that multi-modal integration significantly outperforms unimodal approaches supports resource allocation toward data fusion capabilities rather than isolated prediction systems.

For **academic literature**, the research advances understanding of how predictive analytics and prescriptive optimization can be meaningfully integrated. The framework bridges machine learning and operations research traditions that have developed largely in isolation.

For **future researchers**, the validated framework provides a foundation for extension to other hazard types, infrastructure systems, and geographic contexts, while the identified limitations delineate specific areas requiring further investigation.

1.6 Scope and Limitations

This study focuses on **flood-related extreme weather events** affecting **urban transportation and power infrastructure** in **three mid-sized metropolitan areas** across different climate zones. The temporal scope encompasses the **hazard onset through response phase (first 24 hours)** of flood events.

Exclusions: The framework does not address earthquake, wildfire, or multi-hazard scenarios. It does not model long-term recovery or reconstruction phases. Rural infrastructure and agricultural systems are outside scope.

Limitations: Some infrastructure sensor data were simulated due to availability constraints, introducing uncertainty into validation results. The framework assumes historical patterns retain predictive validity, which may not hold under accelerating climate change. Organizational implementation barriers were assessed qualitatively rather than through controlled deployment.

2. Literature Review

2.1 Conceptual Review

Multi-Modal Decision Support. Multi-modal decision support refers to computational systems that integrate heterogeneous data types—imagery, time-series sensor readings, text, and geospatial data—to inform decision-making processes. The conceptual foundation rests on the principle that different data modalities capture complementary information about complex phenomena, and that fusion of these sources yields superior inference compared to any single modality [3][13]. In emergency management contexts, multi-modal approaches have demonstrated improved detection accuracy for disaster events by combining satellite imagery with social media signals and meteorological observations [18].

Urban Infrastructure Vulnerability. Vulnerability denotes the propensity of infrastructure systems to experience harm or disruption when exposed to hazard events. It comprises three components: exposure (presence in hazard zones), sensitivity (degree of susceptibility to damage), and adaptive capacity (ability to maintain function under stress) [15]. Infrastructure vulnerability is inherently multi-dimensional, varying across physical, functional, and social domains. Assessment approaches range from physically-based fragility curves to indicator-based composite indices, with recent work increasingly employing machine learning for data-driven vulnerability mapping [10].

Emergency Resource Allocation. This construct encompasses the decision processes and optimization methods by which scarce resources—personnel, equipment, supplies, and facilities—are distributed to affected areas during emergencies. Traditional allocation relies on pre-established plans derived from worst-case scenario planning [7]. Dynamic allocation methods adapt to evolving conditions using real-time data, but require predictive inputs to function effectively. The fundamental tension involves balancing efficiency (minimizing cost or maximizing coverage) against equity (ensuring fair treatment across affected populations) [17].

2.2 Theoretical Framework

This study draws on three theoretical perspectives that together explain how decision-support systems can improve emergency outcomes.

Prospect Theory [Kahneman & Tversky] provides the decision-theoretic foundation. Emergency managers operate under conditions of risk and uncertainty where standard expected-utility maximization often fails to describe observed behavior. Prospect theory explains why decision-makers exhibit loss aversion (weighting potential losses more heavily than equivalent gains) and reference-dependent preferences. This has direct implications for resource allocation: managers may over-allocate resources to protect against catastrophic outcomes while under-investing in efficient, lower-risk interventions. The framework's optimization module incorporates these behavioral insights by explicitly modeling both efficiency and equity objectives rather than assuming purely rational allocation.

Situational Awareness Theory [Endsley] explains the cognitive and informational requirements for effective emergency response. Effective decision-making during crises requires three hierarchical levels: perception of critical data elements, comprehension of their meaning, and projection of future states. Traditional emergency management systems often fail at the projection level, providing current-status information without actionable forecasts. The multi-modal framework addresses this by integrating predictive modeling that directly supports projection-level situational awareness [16].

Multi-Criteria Decision Analysis (MCDA) provides the methodological bridge between prediction and allocation. MCDA frameworks enable systematic evaluation of alternatives against multiple, often conflicting criteria—a condition endemic to emergency resource allocation [19]. The proposed framework operationalizes MCDA principles through a stochastic optimization formulation that weighs vulnerability severity, population impact, allocation cost, and equity considerations.

2.3 Empirical Review

Bhuiyan et al. (2026) [1] developed CrisisAI-MIS, a unified multi-modal intelligence framework for real-time crisis prediction and management. The system integrated IoT sensor data, cloud analytics, and AI-driven prediction to address fragmented data flows and delayed detection in organizational crisis response. While demonstrating improved predictive capability compared to baseline classifiers (Logistic Regression, Decision Trees, SVM), the framework focused primarily on organizational crises (system failures, cyberattacks) rather than natural hazard-induced infrastructure failures. Additionally, the study did not incorporate spatial optimization for resource allocation or validate against real-time emergency response scenarios.

Situe et al. (2025) [10][20] proposed an attention-based deep learning framework for urban flood damage and risk assessment. The LSTM-SegNet-MSA hazard prediction model achieved high accuracy in predicting flood depth and extent, while an ES-DeepLab segmentation model characterized land use vulnerability. The framework calculated expected annual damage at pixel-level resolution. However, the study did not address operational resource allocation, and its vulnerability assessment remained static rather than dynamically updated during hazard events.

Fu et al. (2024) [5] developed a hybrid CNN-RNN framework for urban flood prediction that achieved computational efficiency approximately 125 times faster than traditional process-based models. Bayesian optimization identified seven critical spatial driving factors, and the LSTM-DeepLabv3+ model achieved Nash-Sutcliffe Efficiency of 0.9730 under varied rainfall conditions. While this represented significant advancement in flood prediction speed and accuracy, the output remained hazard prediction rather than infrastructure vulnerability assessment, and no resource allocation component was included.

Tabari (2020) [2] and related climate studies have established the empirical basis for increased extreme precipitation under climate change, providing the hazard context for this research. However, climate projection studies rarely translate their findings into operational decision-support frameworks for municipal emergency management.

Two-stage stochastic optimization approaches [7] have been applied to metro system emergency management under uncertain storm floods, demonstrating that dynamic resource allocation during hazard events outperforms pre-event static allocation. This work established the optimization foundation for dynamic allocation but relied on simplified vulnerability assumptions rather than multi-modal prediction inputs.

Multimodal geospatial reasoning [18] using foundation models has demonstrated that combining social media, news, and weather data outperforms traditional hotspot detection for disaster identification. The approach achieved superior precision and recall with only ten content items per prediction. However, this work focused on detection rather than the subsequent resource allocation decision problem.

2.4 Research Gap

The literature review reveals a consistent pattern: predictive modeling, vulnerability assessment, and resource optimization have developed as largely independent research streams. **No validated framework exists that integrates multi-modal vulnerability prediction with dynamic resource allocation optimization in a form operationally deployable for real-time emergency decision support.** Existing prediction models [1][5][10] produce hazard forecasts but not actionable allocation recommendations. Optimization studies [7][17] assume vulnerability inputs that current prediction systems cannot reliably provide at operational timescales. Multi-modal fusion techniques [3][13][18] have not been tested in integrated decision-support architectures.

This study fills this gap by developing and validating an integrated framework that (1) fuses multi-modal data streams for vulnerability prediction, (2) translates predictions into resource allocation recommendations through stochastic optimization, and (3) demonstrates operational performance against historical flood events. The framework explicitly addresses the integration challenge that has limited prior work to either prediction or optimization in isolation.

3. Methodology

3.1 Research Design

This study employs a **quantitative, design-based research design** combining retrospective model development with prospective simulation validation. The design proceeds in three phases: (1) development and training of the multi-modal vulnerability prediction model using historical data, (2) integration with a dynamic resource allocation optimization module, and (3) validation through counterfactual simulation against historical flood events.

This design is appropriate because the research question requires both predictive model development and evaluation of decision-support performance under controlled conditions where ground truth is available. Design-based research permits iterative refinement of the framework while maintaining systematic evaluation protocols.

3.2 Study Area / Population

The framework is developed and validated using data from three urban regions selected for diversity in climate, infrastructure characteristics, and flood exposure:

1. **Region A (Mid-Atlantic U.S.):** Dense urban core with aging combined sewer infrastructure, exposed to hurricane-related flooding.
2. **Region B (Southeast Asia):** Rapidly urbanizing metropolitan area with monsoon-driven pluvial flooding and mixed infrastructure vintages.
3. **Region C (Western Europe):** Medium-density city with modern drainage systems, exposed to riverine flooding from extreme precipitation.

These regions provide variation in infrastructure typology, hazard characteristics, and data availability, supporting assessment of framework generalizability.

3.3 Sample Size and Sampling Technique

The unit of analysis is the **infrastructure node**—defined as critical facilities (substations, water treatment plants), transportation segments (bridges, major intersections), and communication hubs. A total of **847 infrastructure nodes** were identified across the three regions (Region A: 312; Region B: 298; Region C: 237).

Stratified random sampling was employed to select training and validation subsets, with stratification by infrastructure type and historical flood exposure. The training set comprised 70% of nodes ($n=593$) and the validation set 30% ($n=254$). This stratification ensures representation of all infrastructure categories and exposure levels in both training and validation.

3.4 Data Collection Methods

Data were collected from multiple sources spanning 2018-2024:

Meteorological data: Precipitation forecasts and observations from regional weather networks, including rainfall intensity, duration, and spatial distribution at 1-hour temporal resolution.

Infrastructure sensor data: Water level sensors, power system SCADA readings, and traffic monitoring data from municipal utility and transportation agencies. Where sensor data were incomplete, values were imputed using spatial interpolation and temporal gap-filling algorithms.

Geospatial indicators: Elevation models, impervious surface coverage, drainage network characteristics, and building footprint data from open geospatial databases and municipal GIS systems.

Historical flood impact records: Documented infrastructure disruptions (power outages, road closures, water service interruptions) from emergency management after-action reports.

Data were partitioned into training (2018-2022) and validation (2023-2024) periods to assess temporal generalization.

3.5 Research Instruments

The framework was implemented using Python 3.11 with the following libraries: TensorFlow/Keras for neural network components, XGBoost for gradient boosting classification, PyTorch Geometric for spatial graph operations, and Gurobi for stochastic optimization.

Preprocessing pipeline: Data normalization (z-score for continuous variables, one-hot encoding for categorical), temporal alignment to common 1-hour grid, and spatial aggregation to infrastructure node catchments. Missing data were handled through multivariate imputation by chained equations (MICE) for sensor gaps under 6 hours; longer gaps triggered exclusion.

Multi-modal fusion architecture: The prediction component employs a late-fusion design in which modality-specific encoders process meteorological sequences, sensor time-series, and spatial vulnerability features independently before concatenation and joint classification. This architecture permits modality-specific preprocessing and handles heterogeneous sampling rates naturally [3].

3.6 Validity and Reliability

Content validity: Feature selection was guided by domain literature on flood vulnerability factors and reviewed by three emergency management practitioners and two infrastructure engineers. Features retained only if judged theoretically relevant and operationally obtainable during events.

Predictive validity: Model performance assessed on temporally held-out validation data (2023-2024 events) not used in training. Performance metrics include accuracy, precision, recall, F1-score, and area under ROC curve (AUC).

Inter-rater reliability: Historical disruption records were independently coded by two research assistants; Cohen's $\kappa = 0.87$ indicated substantial agreement. Discrepancies resolved through consensus review.

3.7 Data Analysis Techniques

Vulnerability prediction models compared:

1. **Unimodal baselines:** Logistic regression using meteorological features only; random forest using sensor features only.
2. **Multi-modal machine learning:** XGBoost with concatenated features; LSTM network with sequential meteorological input.
3. **Proposed hybrid architecture:** LSTM encoder for temporal meteorological sequences, feedforward network for sensor features, gradient boosting classifier for spatial vulnerability features, with late fusion of modality embeddings.

Model performance evaluated using **5-fold temporal cross-validation** (expanding window) and final assessment on the 2023-2024 holdout period.

Resource allocation optimization: A two-stage stochastic program allocates resources (personnel teams, pumps, generators) under demand uncertainty characterized by the vulnerability prediction distribution. The objective function maximizes expected population coverage subject to equity constraints (no region receives less than 70% of proportional allocation). Solution obtained via scenario-based decomposition [7].

Comparison benchmarks: (a) static allocation based on historical average vulnerability; (b) allocation proportional to population only; (c) allocation based on single-modality prediction (meteorological only).

3.8 Ethical Considerations

This study used **de-identified, publicly available data** from municipal open data portals and federal/regional emergency management databases. No personally identifiable information (PII) or protected health information (PHI) was accessed. Infrastructure disruption records contained no individual-level data. The research protocol was reviewed by the university Institutional Review Board and determined to be **exempt from human subjects review** under 45 CFR 46.104(d)(4) as analysis of existing, de-identified data.

4. Results

4.1 Data Presentation

Table 1. Infrastructure Node Characteristics by Region (2018-2024)

Indicator	Region A (n=312)	Region B (n=298)	Region C (n=237)
Mean elevation (m)	18.4 (12.7)	7.2 (5.1)	42.6 (28.3)
Impervious surface (%)	76.3 (14.2)	62.8 (19.7)	54.1 (16.8)
Historical flood events	23	31	14
Nodes with sensor coverage	68.4%	41.2%	89.7%
Mean annual disruptions	2.3 (1.8)	4.1 (2.7)	1.2 (0.9)

Table 1 presents regional variation in infrastructure and hazard exposure. Region B exhibits highest disruption frequency, consistent with monsoon-driven flooding and lower sensor coverage.

Table 2. Model Performance Comparison (Validation Period)

Model	Accuracy	Precision	Recall	F1	AUC
Meteorological-only (LR)	0.712	0.689	0.724	0.706	0.758
Sensor-only (RF)	0.734	0.712	0.751	0.731	0.779
XGBoost (concatenated)	0.841	0.829	0.856	0.842	0.891
LSTM (temporal)	0.823	0.805	0.847	0.825	0.876
Proposed Hybrid	0.894	0.882	0.906	0.894	0.938

Table 2 shows the proposed hybrid architecture achieves highest performance across all metrics, with 89.4% accuracy and 0.938 AUC.

4.2 Analysis of Results

The hybrid multi-modal framework significantly outperforms unimodal and single-architecture baselines ($p < 0.001$, McNemar's test). Improvement over best single-modality baseline (sensor-only RF) is 16.0 percentage points in accuracy and 0.159 in AUC.

Feature importance analysis (via SHAP values) identified the following top predictors of infrastructure disruption:

1. **Antecedent precipitation index (24-hour)** — weight 0.187
2. **Local elevation relative to drainage outlet** — weight 0.163
3. **Sensor-detected water level rise rate** — weight 0.141
4. **Impervious surface connectivity** — weight 0.118
5. **Power system load volatility** — weight 0.097

Temporal features (precipitation intensity patterns) contributed 34% of predictive power; spatial features (elevation, imperviousness) contributed 41%; sensor features (real-time water levels, system states) contributed 25%.

Resource allocation performance: The dynamic allocation module achieved 42% reduction in allocation latency compared to static pre-positioning (mean response time: 2.8 hours vs. 4.8 hours, $p < 0.01$). Coverage equity improved by 28% as measured by Gini coefficient reduction (0.34 to 0.245). Resource utilization efficiency increased by 31% compared to population-proportional allocation.

5. Discussion

5.1 Interpretation

The finding that multi-modal integration achieves 89.4% accuracy addresses Research Question 1 by demonstrating that complementary information across meteorological, sensor, and geospatial domains is essential for reliable vulnerability prediction. This aligns with prior work on multi-modal fusion in disaster detection [3][18] and extends it to infrastructure-specific vulnerability assessment. The feature importance analysis reveals that temporal meteorological patterns and spatial vulnerability characteristics contribute nearly equally to predictive power, suggesting that neither dimension alone is sufficient.

The comparison against baseline methods addresses Research Question 2. The proposed framework's 42% reduction in allocation latency represents operationally significant improvement

over static methods, confirming theoretical arguments that dynamic allocation outperforms pre-event planning under uncertain conditions [7]. The improvement in coverage equity (28% Gini reduction) demonstrates that efficiency and equity objectives are not necessarily conflicting when vulnerability predictions are accurate—a finding that challenges assumptions in some prior allocation literature [17].

The hybrid architecture's performance advantage over XGBoost with concatenated features (89.4% vs. 84.1% accuracy) indicates that the late-fusion design captures modality-specific structures that simple feature concatenation obscures. This supports the technical rationale for modality-specific encoders and aligns with architectural principles from multi-modal learning literature [3].

Regarding Research Question 3, qualitative assessment identified three primary implementation barriers: (1) data integration challenges where municipal agencies maintain siloed data systems, (2) computational infrastructure requirements exceeding current emergency operations center capabilities, and (3) organizational resistance to algorithmic decision recommendations. These barriers are addressable through phased implementation and training but require explicit attention in deployment planning.

5.2 Implications

Academic implications: This study extends multi-modal learning theory by demonstrating that late-fusion architectures with modality-specific encoders outperform feature concatenation for spatially heterogeneous prediction tasks. The finding that temporal and spatial features contribute nearly equally challenges the common assumption that real-time sensor data dominates predictive value in operational contexts. The integration of prediction and optimization within a single decision-support framework advances understanding of how machine learning and operations research methods can be meaningfully coupled.

Practical implications: Emergency management agencies should prioritize investment in multi-modal data fusion capabilities rather than isolated prediction improvements. The framework provides specific guidance: vulnerability predictions at 89.4% accuracy with 6-hour lead time can support dynamic resource allocation decisions. Administrators should monitor prediction precision and recall metrics, retraining models when performance degrades below 85% accuracy. Expected lead time for allocation decisions is 2-3 hours from prediction availability, enabling responsive adaptation to evolving hazard conditions.

5.3 Limitations

1. **Simulated sensor data:** Approximately 31% of sensor records in Region B were simulated due to coverage gaps, introducing potential bias toward performance characteristics of regions with better sensor networks.

2. **Historical pattern stability assumption:** The framework assumes relationships between predictors and outcomes remain stable, an assumption threatened by accelerating climate change and infrastructure modifications.
3. **Limited hazard scope:** Validation focused on flood events; performance for other extreme weather types (hurricanes, winter storms) is untested.
4. **Organizational barriers not empirically tested:** Implementation challenges were assessed qualitatively; controlled deployment studies are needed to validate practical feasibility.

5.4 Future Research Directions

1. **Extension to other hazard types and infrastructure systems:** Testing the framework for wildfire, earthquake, and multi-hazard scenarios, and extending to water, communication, and healthcare infrastructure.
2. **Longitudinal deployment studies:** Partnering with municipal emergency management agencies for prospective validation during actual events, enabling assessment of organizational adoption and decision quality.
3. **Climate non-stationarity adaptation:** Developing methods to detect and adapt to distributional shifts in hazard-vulnerability relationships under climate change.
4. **Human-AI teaming optimization:** Investigating how emergency managers interact with and calibrate trust in algorithmic decision recommendations under high-stress conditions.

6. Conclusion

This study developed and validated a multi-modal decision-support framework achieving 89.4% accuracy in urban infrastructure vulnerability prediction and 42% reduction in emergency resource allocation latency. The integrated architecture—combining LSTM temporal encoding, gradient boosting spatial classification, and stochastic optimization—demonstrates that meaningful integration of prediction and prescriptive analytics is technically feasible and operationally impactful. The framework provides a replicable model for municipalities seeking to strengthen resilience through data-driven decision support. For emergency management practitioners, the key takeaway is that multi-modal integration substantially outperforms isolated prediction or static allocation approaches, justifying investment in data fusion infrastructure. As extreme weather intensifies, such integrated decision-support systems will transition from innovation to necessity for urban resilience.

References

1. Bhuiyan, M. S., Hasan, R., Ahmed, M. J., Hasan, M., Ahmed, K. R., & Islam, R. (2026, April). CrisisAI-MIS: A unified multi-modal intelligence framework for real-time crisis prediction and management. In *2026 IEEE 2nd International Conference on Quantum Photonics, Artificial Intelligence & Networking (QPAIN)* (pp. 1-6). IEEE.
2. Tabari, H. (2020). Climate change impact on flood and extreme precipitation increases with water availability. *Scientific Reports*, *10*(1), 13768. <https://doi.org/10.1038/s41598-020-70816-2>
3. AI-powered multimodal disaster prediction and response system using remote sensing and edge analytics. (2026). In *2025 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS)* (pp. 1-8). IEEE. <https://doi.org/10.1109/ICCCIS61093.2025.11362424>
4. Vitoriano, B., Ortuño, M. T., Tirado, G., & Montero, J. (2013). A multi-criteria optimization model for humanitarian aid distribution. *Journal of Global Optimization*, *51*(2), 189-208.
5. Fu, G., Jin, T., & Teng, S. (2024). A novel hybrid deep learning framework for urban flood prediction with spatiotemporal feature fusion. *Journal of Hydrology*, *630*, 130685. <https://doi.org/10.1016/j.jhydrol.2024.130685>
6. Saksrisathaporn, K. (2015). *Decision aid models for humanitarian logistics* [Doctoral dissertation, Université Lumière Lyon 2].
7. Two-stage stochastic optimization for emergency management of metro systems under uncertain storm floods. (2025). *Reliability Engineering & System Safety*, *261*, 110826. <https://doi.org/10.1016/j.ress.2025.110826>
8. Chamundeswari, G., Jayendra, D., Thatipudi, G., Begum, A., & Mudalkar, P. (2026). A LSTM–XGBoost integrated model for smart resilience networks. In *Smart Resilience Networks* (pp. 5-1–5-20). Taylor & Francis. <https://doi.org/10.1201/9781003642145-5>
9. Lu, S., & Zhuang, J. (2023). A systematic review of machine learning applications in crisis management. *Journal of Optimization in Industrial Engineering*, *16*(2), 15-40.
10. Situ, Z., Zhong, Q., Zhang, J., Teng, S., Ge, X., Zhou, Q., & Zhao, Z. (2025). Attention-based deep learning framework for urban flood damage and risk assessment with improved flood prediction and land use segmentation. *International Journal of Disaster Risk Reduction*, *116*, 105165. <https://doi.org/10.1016/j.ijdrr.2024.105165>

11. Scopus. (2026). CrisisAI-MIS: A unified multi-modal intelligence framework for real-time crisis prediction and management [Citation record]. *Scopus*. <https://www.scopus.com>
12. Arias, P., Bellouin, N., Coppola, E., Jones, R., Krinner, G., Marotzke, J., Naik, V., Palmer, M., Plattner, G.-K., Rogelj, J., et al. (2021). *Climate change 2021: The physical science basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. IPCC.
13. Multimodal fusion of remote sensing and meteorological data for flood prediction. (2026). In *2026 International Conference on Recent Trends in Computing and Communication (ICRTCC)* (pp. 1-7). IEEE. <https://doi.org/10.1109/ICRTCC63636.2026.11570311>
14. Del Frate, N. (2025). *Multi-criteria decision analysis in public health emergencies: A systematic review* [Master's thesis, Università del Piemonte Orientale].
15. Pourghasemi, H. R., et al. (2026). From prediction to prevention: An explainable GeoAI framework for flood susceptibility and urban exposure assessment using machine and deep learning models. *Sustainable Development*, *34*(1), 1-20. <https://doi.org/10.1002/sd.70772>
16. Bhuiyan, M. S., Hasan, R., Ahmed, M. J., Hasan, M., Ahmed, K. R., & Islam, R. (2026). CrisisAI-MIS: A unified multi-modal intelligence framework for real-time crisis prediction and management. In *2026 IEEE 2nd International Conference on Quantum Photonics, Artificial Intelligence & Networking (QPAIN)*. IEEE. <https://doi.org/10.1109/qpain69676.2026.11545683>
17. Emergency supply distribution and shelter selection optimization under urban flooding. (2026). *Geomatics, Natural Hazards and Risk*, *17*(1), 2719317. <https://doi.org/10.1080/19475705.2026.2719317>
18. Hanny, D., Dastidar, K. G., Wieland, M., et al. (2026). Towards multimodal geospatial reasoning: A foundation model approach for disaster detection from social media, news, and weather data. *Natural Hazards*, *121*(4), 1-45. <https://doi.org/10.1007/s11069-026-08191-w>
19. Multi-criteria decision making in emergency management: A comprehensive review. (2025). *Mathematics*, *13*(5), 731. <https://doi.org/10.3390/math13050731>
20. Situ, Z., et al. (2025). Attention-based deep learning framework for urban flood damage and risk assessment. *International Journal of Disaster Risk Reduction*, *116*, 105165.