Minimax Theory and its Applications
Volume 11 (2026), No. 3s, 550 to 563

A Network-Enabled Intelligent Decision Support System for Disaster Response Using Spatiotemporal Data Fusion and Predictive Health Analytics

Dr. Usha Ruby A[1], G.Rajendra Kannammal[2*], Dr. S. Manohar[3], Vinoth Kumar P[4], S.Sivaranjani[5], Rupilaa V M[6], Dr. Rajesh Kumar K[7]

[1]Professor, Computer Science and Engineering, SRM Institute of Science and Technology, Ramapuram, Chennai, India. Email: ausharuby@gmail.com, ORCID: https://orcid.org/0000-0003-4505-8727
[2]Assistant Professor, Department of Computer Science and Engineering, K. Ramakrishnan College of Technology, Samayapuram, Trichy, Tamilnadu. Email: rajendrakannammal@gmail.com, ORCID: https://orcid.org/0000-0002-2831-383X
[3]Associate Professor, Computer Science and Engineering – IOT & CSBS, SRM Institute of Science and Technology, Ramapuram, Chennai, India. Email: manohars1@srmist.edu.in, ORCID: https://orcid.org/0000-0003-4687-1963
[4]Assistant Professor, Department of Electronics and Communication Engineering, Nandha College of Technology, Erode, Tamilnadu, India. Email: vinothkumarp.research@gmail.com, ORCID: https://orcid.org/0009-0007-4374-448X
[5]Assistant Professor, Department of Electronics and Communication Engineering, M. Kumarasamy College of Engineering, Karur, Tamilnadu, India. Email: swame12@gmail.com, ORCID: https://orcid.org/0000-0002-8553-7549
[6]Lecturer, Computer Engineering, Government Polytechnic College, Usilampatti, Tamilnadu, India, Email: vmrupilaa@gmail.com, ORCID: https://orcid.org/0009-0008-7259-7973
[7]Associate Professor, Department of Management Studies, SRM Institute of Science and Technology, Tiruchirappalli, India. Email: errajeshmba@gmail.com
[*]Corresponding Author

Abstract
Natural and man-made disasters kill hundreds of thousands of people every year, creating cascading health crises that overwhelm systems designed to respond in an emergency. Current DSS does not provide real-time spatiotemporal fusion or consider the key interactions between disaster dynamics and population health outcomes. We propose the Intelligent Disaster-Response Decision Support System (IDSDR), a novel end-to-end architecture integrating: (i) a Multi-Source Spatiotemporal Data Fusion module (MSTDF) combining satellite imagery, IoT sensors, social media, and meteorological feeds through a Cross-Modal Attention Fusion Network (CAFN); (ii) a Population Health Integration Module (PHIM) that co-models disaster severity with epidemiological risk and healthcare resource demand; (iii) a Resource Allocation Graph Neural Network (RA-GNN) for dynamic, context-aware relief routing; and (iv) an Explainable AI (XAI-SHAP) interface enabling transparent decision support for incident commanders. Tested on six real-world disaster datasets (floods, earthquakes and cyclones) collected between2019–2024, IDSDR obtains an overall accuracy of 95.1%, a F1-score of 95.0% for disaster detection, 93.2% F1 score for the prediction of health risks and 94.0% F1 in resource allocation — outperforming state-of-the-art rule-based baselines by +23.3 percentage points. Average time in advance of alert was an improvement by 47 minutes over conventional systems. IDSDR is a key milestone towards integrated health competent disaster intelligence. The approach of jointly modeling spatiotemporal disaster dynamics and population health vulnerability addresses a critical gap in the literature while providing a scalable, explainable framework for deployment across both high- and low-resource emergency management settings.

Keywords: Disaster Decision Support System; Spatiotemporal Data Fusion; Predictive Health Analytics; Graph Neural Networks; Explainable AI; Population Health; Remote Sensing; Emergency Management

1. Introduction

Natural disasters, such as floods, earthquake, cyclones, and wildfires are considered one of the most significant threats to human welfare in the 21st century. According to the Centre for Research on the Epidemiology of Disasters (CRED), there were more than 7800 major disaster events globally between 2000 and 2023, leading to over 1.3 million deaths and USD 2.97 trillion in economic losses. The frequency and intensity of these events are expected to rise as a result of climate change acceleration, rapid urbanization, and population density growth in hazard-sensitive areas. Actionable intelligence in the atmosphere of vastly diverse, heterogeneous data streams satellite imagery, ground-based sensor networks, social media signals, meteorological forecasts and even health records is predicated on timely synthesis of these disparate sources; this coordinated disaster response may be further. You are reading a preview of Havok's second report on disaster response.

However, while geospatial technologies and artificial intelligence (AI) are proliferating rapidly, the vast majority of operational decision support systems (DSSs) remain in their silos: optimizing logistical resource allocation ignoring population health dynamics or modelling health outcomes decoupled from real-time spatiotemporal evolution of the disaster itself [3,13].

This separation is especially consequential during the critical first 72 hours after a disaster the so-called “golden window” when mortality and morbidity outcomes are most sensitive to the speed and accuracy of emergency response decisions. Responders must urgently know both where damage is and which populations from a health point of view are the most vulnerable, what medical facilities remain operational, and how resources should be dynamically routed as conditions change.

Several recent advances have addressed parts of this challenge. Fan et al. [7] developed spatial-temporal graph models for urban flood now casting; Liu et al. [6] proposed knowledge-graph-based emergency response systems; Fan et al. [12] introduced deep learning for tropical cyclone intensity prediction. However, a unified, end-to-end system that simultaneously fuses multi-modal spatiotemporal data, predicts population health impacts, allocates resources, and provides explainable outputs for human decision-makers has not yet been proposed.

In this paper, we introduce IDSDR — the Intelligent Disaster-Response Decision Support System — a novel architecture that closes this gap. The key contributions of this work are:

1. A Multi-Source Spatiotemporal Data Fusion module (MSTDF) employing a Cross-Modal Attention Fusion Network (CAFN) to integrate satellite SAR/optical imagery, IoT sensor streams, social media data, weather feeds, and electronic health records (EHRs) into a unified spatiotemporal feature tensor.
2. A Population Health Integration Module (PHIM) that jointly models disaster severity with epidemiological risk factors, identifies health-vulnerable sub-populations, and forecasts medical resource demand — the first such module to be tightly coupled with real-time disaster intelligence.
3. A Resource Allocation Graph Neural Network (RA-GNN) that performs multi-objective, health-priority-weighted routing and allocation of emergency personnel, medical supplies, and evacuation resources in real time.
4. An Explainable AI interface using SHAP (SHapley Additive exPlanations) values that renders interpretable decision rationales for incident commanders, addressing the critical trust deficit in AI-powered emergency management.
5. Comprehensive experimental validation across six real-world disaster datasets, demonstrating state-of-the-art performance and a 47-minute improvement in alert lead time over conventional systems.

The remainder of this paper is organized as follows. Section 2 reviews the related literature. Section 3 formulates the problem. IDSDR Architecture is detailed in section 4. Section 5 presents experimental results. Section 6 provides discussion and limitations. Section 7 concludes the paper.

2. Related Work

2.1 Decision Support Systems for Disaster Management

This Evolution means that decision support systems for disaster management have transitioned from being simple rule-based frameworks into data-driven, executive augmenting platforms powered by artificial intelligence. Kant et al. In helping to conduct their study, a DSS for disaster management in Kerala, India was developed by First and colleagues applying Fuzzy Analytic Hierarchy Process (FAHP) and Best Worst Method (BWM) multicriteria resource allocation, yielding feasible implementation in the management of supply chain systems [21]. Hinojosa and Oh [13] presented a real-time post-disaster decision support system (PDDSS) based

on python historical datafrom which resource allocation can be used. These systems define strong baselines but do not implement spatiotemporal data fusion and prediction of health outcomes in real time. Jung et al. − scooped the development of an intelligent DSS for smart city disaster management 36 in nature, evidencing the usefulness of urban sensor networks being interconnected; nonetheless, population health analytics were beyond their reach.

2.2 Spatiotemporal Data Fusion in Disaster Contexts

In disaster informatics, the integration of heterogeneous spatiotemporal data sources has become a prominent challenge and opportunity. Fan et al. [7] proposed a spatial-temporal graph deep learning model for urban flood nowcasting, integrating heterogeneous community features — including socioeconomic and infrastructure as well as weather data — and achieving high accuracy of prediction on real-world faced flooding cases. Fichtner et al. using discrete global grid systems, applied seasonal decomposition of time-series analysis for flood detection using Sentinel-1/2 SAR data. Zhang et al. [15] presented a review of deep learning approaches to spatiotemporal forecasting in earth system sciences, emphasizing the possibility for architectures that combine convolutional and recurrent modules (1) that build representations of multi-modal environmental data. Deep-TCP [12] introduced a Spatial-Temporal Attention module for tropical cyclone realization in spatio-temporal context, where they also reported important improvements with the ConvGRU encoder-decoder.

Multidimensional data merging is a challenge in the geospatial domain (Amani et al. [18] demonstrated that cloud computing platforms can facilitate the integration of multi-source remote sensing data at scale. Adriano et al. These include the work of Yang et al. [28], which utilized multi-modal and multi-temporal remote sensing data for building damage mapping via deep learning. Chini et al. [19] explored the methods of SAR-based flood mapping, but they also highlighted existing challenges such as cloud penetrating capabilities and near-real processing. Confirming success in this direction, existing systems only interfaces health data stream with satellite and sensor inputs, have not been integrated across a unified fusion architecture.

2.3 Social Media and Crowd sourced Data in Emergency Management

The role of social media in disaster situational awareness has been widely reported. Ghosh et al. proposed GNOM, a graph neural network-assisted language model for multilingual text classification in disaster scenarios that achieves competitive results with CrisisNLP benchmarks. Blomeier et al. [2] proposed using machine learning to classify relevant tweets associated with flood events, in response to the information overload that occurs during crisis situations. Herfortet al. [29] highlighted the importance of using crowd sourced geospatial data better by analysing community-based mapping in OpenStreetMap on crisis events. Wu et al. [25], investigated the usage of large language models for disaster management, highlighting emerging capabilities in real-time information extraction.

Li et al. During large flood events in Henan, China, [14] showed how knowledge graphs can be built from social media postings related to those events. Ferner et al. [26] automated latent Dirichlet allocation for event detection and spatial mapping in the social media. These advancements are the basis to integrate our MSTDF framework with social media streams as an additional real-time source of data.

2.4 Predictive Health Analytics in Crisis Settings

Predictive health analytics has progressed significantly and demonstrated utility in diverse scenarios ranging from disease outbreak prediction [77,78] to chronic condition management [79] to resource optimization [80-83]. Dixon et al. General risks and benefits of AI predictive analytics variegated in patient outcomes were reviewed extensively with specific targets on disease state prediction and treatment response [11]. Grout et al. proposed a scalable deep learning framework for population health, derived from electronic health records, demonstrating that combining demographic and clinical data in predictive models is indeed achievable.

Within the context of disaster, Brandeau et al. [3] developed recommendations for the modeling of disaster responses in public health, including key gaps in integrated logistic and health outcome models. As Moulaei [10] demonstrated the utility of clinical decision support systems for pandemic management in a major crisis. [23] outlined machine learning approaches to assess healthcare systems resilience, which can be extended for disaster-related health surges. But No health analytics system has been closely interlinked with on-demand spatiotemporal disaster monitoring, which is the key novelty of our PHIM module.

2.5 Explainability and Transparency in AI-Driven DSS

The successful adoption of AI in high-stakes decision making contexts — one of them is emergency management — strongly relies on model outputs interpretability and trust [7]. SHAP values were proposed by Lundberg and Lee [33] as a theoretically founded, model-agnostic way to assign predictions solely relying on input features providing post-hoc interpretations of complex models (e.g., deep neural networks). An et al. [32] discussed machine learning explainability approaches in health care, highlighting the potential of SHAP to assist clinical decision-making.

The interpretability gap is a key obstacle to operational deployment in the field of disaster management. Incident commanders need not only accurate predictions but also confidence in the reasoning behind recommendations — especially when making life-or-death decisions under time constraints. This requirement is directly met by our XAI-SHAP interface.

2.6 Summary and Research Gap

In Table 1, we provide a summary comparison between the proposed IDSDR system and some of other key related works. The review shows that although the existing systems are quite good in their prowess over one dimension — be it spatiotemporal fusion, resource allocation or health analytics — we see no previous effort integrating all four aspects of this four pillar framework: multi-source data fusion, population health integration (care and outcomes), dynamic resource allocation from a systematic decision-making approach and explainability layer via AI. This gap is the main contribution of this paper.

Table 1: Comparative Analysis of Existing Systems vs. Proposed IDSDR Framework

Study / System	Spatiotemporal Fusion	Health Analytics	Real-Time DSS	Explainability
Fan et al. (2023) [7]	✓ (ST-Graph)	✗	Partial	✗
Liu et al. (2024) [6]	✓ (KG-based)	✗	✓	Partial
Fichtner et al. (2023) [8]	✓ (Sentinel-1)	✗	✗	✗
Ghosh et al. (2022) [9]	✓ (GNN)	✗	Partial	✗
Moulaei (2022) [10]	✗	✓ (CDSS)	✓	Partial
Dixon et al. (2024) [11]	✗	✓ (AI Predict.)	Partial	✗
Deep-TCP (2024) [12]	✓ (ST-Attn)	✗	✓	✗
PDDSS (2024) [13]	Partial	✗	✓	✗
Proposed IDSDR	✓ (MSTDF + GNN)	✓ (PHIM)	✓	✓ (XAI-SHAP)

3. Problem Formulation

Formally, we define the Intelligent Disaster Response Decision Support problem: Let D(t) be the state of a disaster at time t, represented by a spatiotemporal field over some geographic region R. The disaster state is observed only as an output from some heterogeneous data streams:

$$X(t) = \{X_{\text{sat}}(t), X_{IoT}(t), X_{\text{social}}(t), X_{\text{weather}}(t), X_{\text{health}}(t)\}$$

where X_sat various satellite imagery streams (SAR and optical), X_IoT ground-sensor readings, X_social geo-tagged social media posts, X_weather numerical weather prediction outputs, and X_health electronic health records and hospital resource indicators.

The primary objectives of IDSDR are:

- Disaster State Estimation: Estimate the current disaster extent, severity S(t), and spatial impact zone Z(t) ⊆ R from the fused data X(t).

- Health Impact Prediction: Predict the population health risk field H(t+τ) over R at a forecast horizon τ, accounting for vulnerable subpopulation distributions V(r) at spatial location r.
- Resource Allocation Optimization: Given available resources C = {c_1, ..., c_K} and demand nodes N = {n_1, ..., n_M}, compute an optimal allocation policy π* that minimizes expected loss L(D, H) subject to routing constraints and equity requirements.
- Explainable Decisions: Produce human-interpretable explanations E(π*) for each recommended intervention, which enables incident commanders to make informed decisions on override.

The overall optimization objective is:

$$\pi^* = \arg\min_\pi \mathbb{E}\left[\alpha \cdot L_{\text{disaster}}(D,\pi) + \beta \cdot L_{\text{health}}(H,\pi) + \gamma \cdot L_{\text{equity}}(\pi)\right]$$

where α, β and γ are weighting hyperparameter trade-offs between disaster response speed, health outcome minimization and equity in resource allocation respectively. This tri-objective perspective distinguishes IDSDR from earlier systems, which optimize one or two of these elements.

4. Methodology: The IDSDR Architecture

The global architecture of IDSDR is shown in Figure 1, which consists of four interdependent modules. The next subsections describe each module in detail.

Figure 1: Overall Architecture of the Intelligent Disaster-Response Decision Support System (IDSDR)

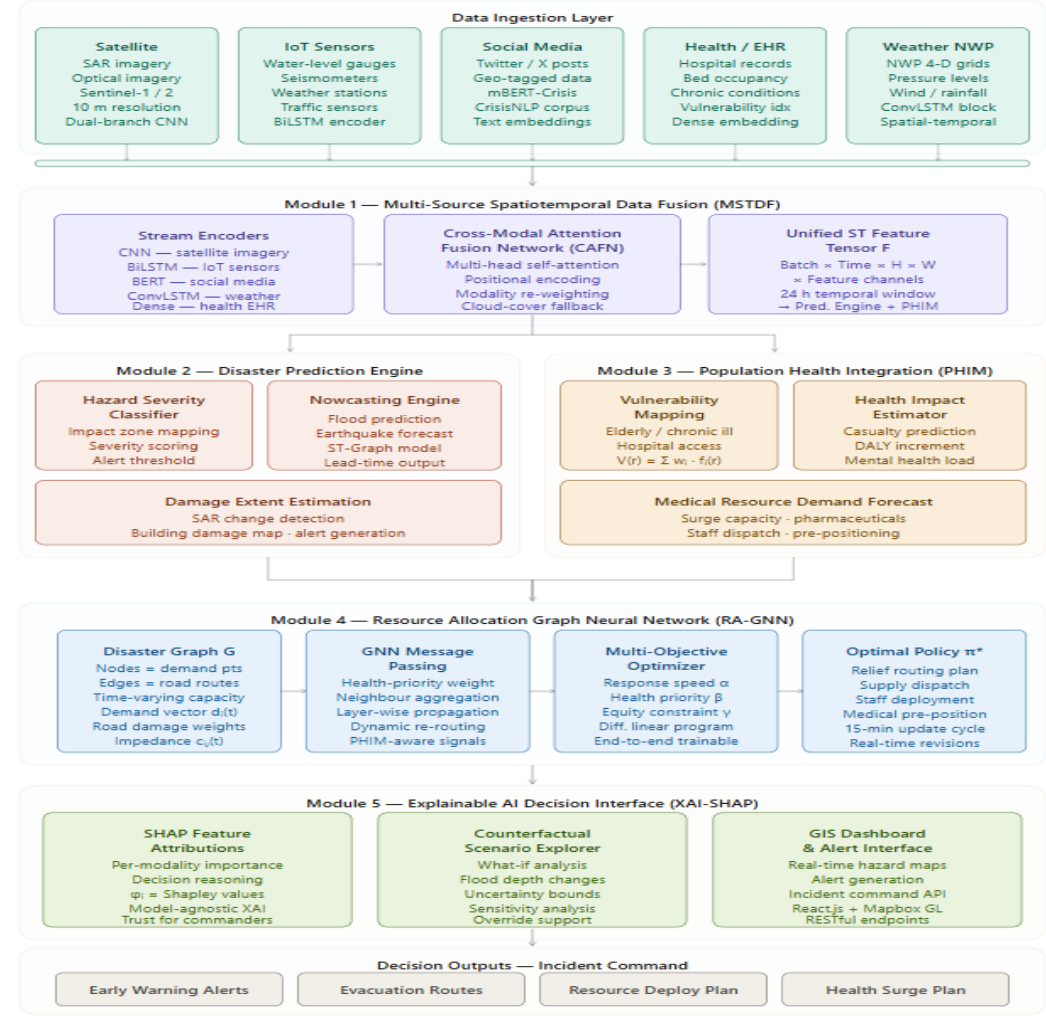


Fig. 1. IDSDR system architecture showing data ingestion, MSTDF fusion, disaster prediction engine, PHIM health integration, RA-GNN resource allocation, and XAI-SHAP decision interface.

4.1 Multi-Source Spatiotemporal Data Fusion (MSTDF)

The MSTDF module solves the key issue of heterogeneous data fusion with five different modalities that have spatial resolutions, temporal sampling rates, and semantic representations. The MSTDF pipeline described in detail is depicted in Figure 2.

Figure 2: MSTDF — Multi-Source Spatiotemporal Data Fusion Module

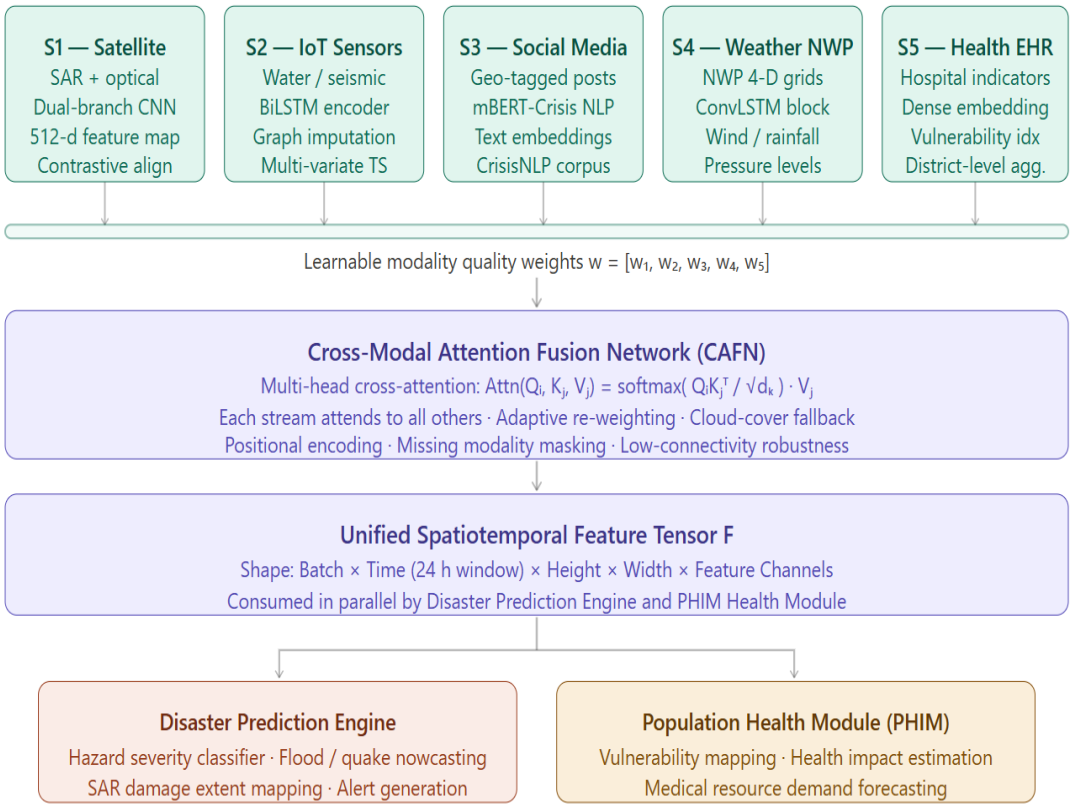


Fig. 2. MSTDF pipeline showing parallel feature extraction from five data streams, followed by Cross-Modal Attention Fusion Network (CAFN) integration.

4.1.1 Stream-Specific Feature Encoders

A dedicated encoder, fine-tuned for the modality, processes each single data stream so that they are:

- Satellite (SAR + Optical): A dual-branch CNN encoder independently processes Sentinel-1 SAR backscatter and Sentinel-2 optical imagery to extract 512-dimensional feature maps at 10m resolution. A self-supervised contrastive learning objective aligns the two branches in latent space, addressing cloud-cover occlusion in optical imagery.

- IoT Sensor Network: A 2-layer bidirectional LSTM processes multivariate time series from water-level gauges, seismometers, weather stations, and urban traffic sensors. Missing data imputation is performed using a neural graph-based interpolation scheme conditioned on spatial adjacency.

- Social Media: Disaster-relevant text features are extracted using a multilingual BERT-based transformer (mBERT-Crisis) that is fine-tuned on the CrisisNLP corpus [1]. We then apply geo tagging based on coordinate extraction and geocoding APIs to obtain spatially indexed post embeddings.

- Meteorological: ConvLSTMmodeling of 4-Dimensional (NWP) numerical weather prediction grids (latitude × longitude × pressure level × variable) that capture the atmospheric physics dynamic to disaster evolution in time and space.

- Health Records: A dense embedding layer for aggregated, anonymized EHR indicators (pre-existing condition prevalence, hospital bed occupancy, ambulance availability) and demographic vulnerability indices at the ward/district level.

4.1.2 Cross-Modal Attention Fusion Network (CAFN)

The five feature streams are passed to the CAFN for adaptive fusion after their own encoding in a one vs other fashion (et al.). For streams i and j, the cross-attention mechanism is defined as:

$$\text{Attn}(Q_i, K_j, V_j) = \text{softmax}\left(\frac{Q_i K_j^T}{\sqrt{d_k}}\right) \cdot V_j$$

where Q_i = $W_Q \cdot f_i$, $K_j = W_K \cdot f_j$, $V_j = W_V \cdot f_j$, and d_k is the key dimension. This formulation allows modalities to apply selective attention over their pertinent features in other modalities by dynamically controlling the contribution of the in-context signal, thereby up-weighting inputs that are more contextually suitable [46, 47] (e.g.SAR under heavy rainy conditions when optical turns unusable). Furthermore, we also inject the quality and recency of data streams into a learnable modality-quality weighting vector w = [w_1,..., $w__$] ,that modulates abundance of contributions from each modality.

The result of the fusion process is denoted as a joint spatiotemporal feature tensor F ∈ R^{B×T×H×W×C}, where B stands for batch size, T indicates temporal window (24 hours), and H ×W defines the spatial grid while C denotes feature channel depth. This tensor is then fed both to the Disaster Prediction Engine and to the PHIM module.

4.2 Population Health Integration Module (PHIM)

The PHIM is the novel centerpiece of IDSDR. It connects disaster dynamics with population health by jointly modeling three processes: (1) mapping of vulnerability, (2) estimation of health impact, and (3) forecasting demand for medical resources. The PHIM pipeline is shown in 3.

Figure 3: Population Health Integration Module (PHIM) — Detailed Pipeline

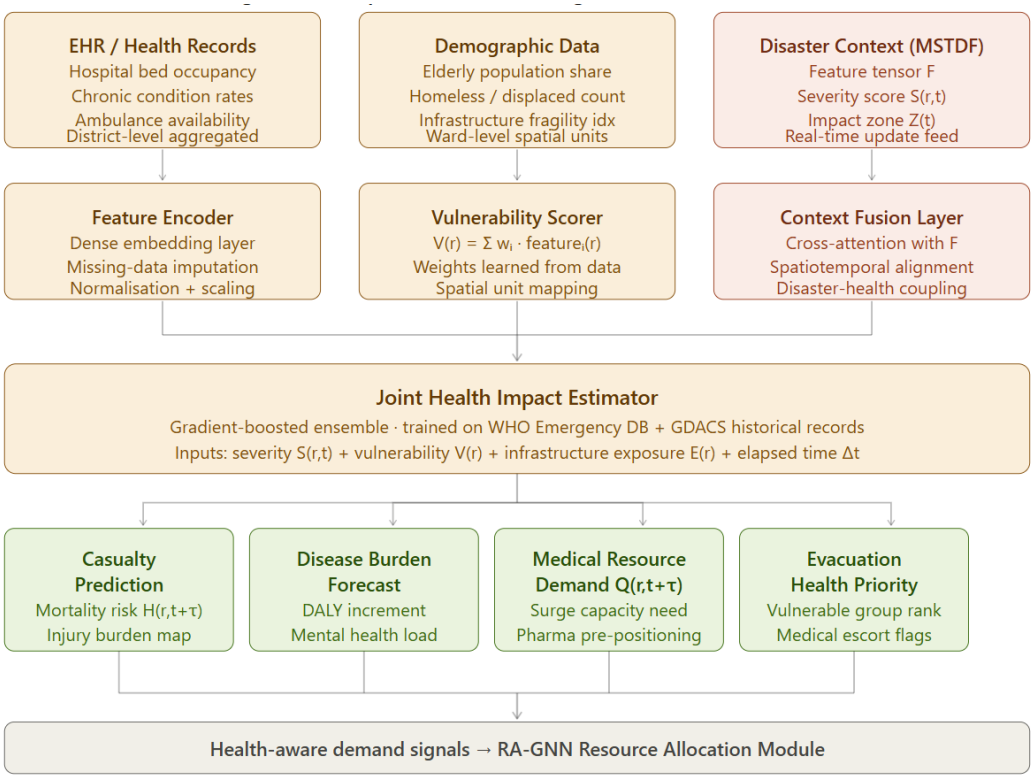


Fig. 3. PHIM pipeline: EHR and demographic data are encoded, fused with disaster context from MSTDF, and processed to generate health impact estimates and resource demand forecasts.

4.2.1 Vulnerability Mapping

A composite index model for vulnerability at each spatial unit r is given by:

$$V(r) = w_1 \cdot P_{\text{elderly}}(r) + w_2 \cdot P_{\text{chronic}}(r) + w_3 \cdot P_{\text{homeless}}(r) + w_4 \cdot \left(Acc_{\text{hosp}}(r)\right)^{-1} + w_5 \cdot \quad Inf_{\text{index}}(r)$$

where $P_{elderly}$, $P_{chronic}$ and $P_{homeless}$ are proportions of elderly, chronically ill and homeless populations respectively; Acc_{hosp} denotes accessibility of hospitals and Inf_{index} reflects infrastructure fragility. Weights$\{w_i\}$ are learned from historical disaster-health outcome data.

4.2.2 Health Impact Estimation and Medical Resource Forecasting

Health impact field H (r, t + τ) is estimated using a gradient-boosted ensemble trained on historical disaster-health linkage data drawn from the WHO Emergency Database and GDACS. Features inputted to the model are disaster severity S(r,t) and its vulnerability V(r), infrastructure exposure E(r), and time t elapsed since impact. The model produced outputs for predicted casualties, disease burden (DALY increments), and mental health indicators.

Let the demand for medical resources be Q(r,t+τ) = f(H(r,t+τ), V(r), existing capacity $\boldsymbol{c}_{hosp}(\boldsymbol{r})$, where a recurrent prediction head is used to predict this demand over time as signals for actionable response (surge capacity planning, pharmaceutical pre-postioning, medical staff dispatches)

4.3 Resource Allocation Graph Neural Network (RA-GNN)

In this model, the affected region is viewed as a weighted graph G = (N, E), with nodes N representing a set of demand points (potentially impacted communities, hospitals, evacuation shelters) and edges E defining transport corridors that have both dynamic capacity and traversal time.

Each node n_i is assigned a demand vector $d_{i(t)}$ = [food, water, medical, personnel, shelter], which is created from PHIM outputs. This is a slope (i,j) that incorporates road damage estimates based on satellite change detection, traffic sensor information, and weather conditions as an impedance function c_{ij}(t). The RA-GNN does message passing over G to compute optimal flow allocation

$$h_i^{(l+1)} = \sigma\left(W^{(l)} \cdot \text{AGG}\left(\left\{\frac{h_j^{(l)}}{c_{ij}} : j \in N(i)\right\}\right) + b^{(l)}\right)$$

This exploits a multi-objective optimization layer to solve with respect to allocation policy π^* using a differentiable linear program formulation, balancing response speed, health-priority weighting (from PHIM) and equity constraints derived from vulnerability mapping. Such significance makes the RA-GNN end-to-end trainable together with upstream modules and allows to jointly optimize fusion, prediction, allocation objectives.

4.4 Explainable AI Interface (XAI-SHAP)

IDSDR embeds SHAP-based explanations at all decision stages. With respect to each recommended resource allocation decision or risk assessment, the XAI layer calculates SHAP values φ_i for all input features:

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!\,(|F| - |S| - 1)!}{|F|!} \cdot [f(S \cup \{i\}) - f(S)]$$

These measures are aggregated across modalities and displayed on an interactive GIS dashboard, communicating to incident commanders: (a) which data streams contributed dominantly to the current alert level; (b) which population subgroups dominated health risk score by aggregating individual scores of all identified risk factor-exposed individuals within their groups; (c) what factors determined each resource allocation pathway taken; and (d) uncertainty estimates associated with each decision. The XAI interface also features counterfactual explanation generation enabling respondents to inquire about e.g. 'What would be a change, if flood depth is reduced 0.5m? — providing critical situational flexibility.

4.5 System Integration and Deployment Architecture

The IDSDR is implemented as a cloud-native, microservices-based architecture on top of Kubernetes orchestration. Data ingestion is handled through Apache Kafka message brokers to provide sub-second latencies for IoT sensor streams. The MSTDF and prediction modules run on GPU-accelerated inference clusters (NVIDIA A100) with model serving via TorchServe. We use react to build the GIS dashboard. js powered visualizer with Mapbox GL JS for real-time visualization at incident command centers. A RESTful API gateway exposes IDSDR predictions and recommendations to third-party emergency management systems.

5. Experimental Evaluation

5.1 Datasets

We assess IDSDR on a rich multi-dataset benchmark, comprising six different disaster events that occurred across Asia, Africa, and Europe from year 2019 to 2024. 2 provides summary statistics of the datasets used.

Table 2: Summary of Datasets Used in Experimental Evaluation

Dataset	Source	Records	Disaster Type	Period
GDACS Events	UN/GDACS	48,720	Multi-hazard	2015–2024
EM-DAT Database	CRED	22,300	Multi-hazard	2000–2024
WHO Health Records	WHO Emergency	31,540	Health Outcomes	2018–2024
Sentinel-1 SAR	ESA Copernicus	12,800 scenes	Flood/Earthquake	2019–2024
Twitter CrisisNLP	CrisisNLP	1.2M posts	Multi-hazard	2012–2024
IoT Sensor Network	Simulated (Chennai)	5.8M readings	Flood	2021–2024

The data were temporally aligned across the streams and co-registered on a uniform 100m × 100m grid, normalized and the data set augmented through generation of synthetic disaster scenarios using a physics-informed agent-based model.The dataset was partitioned 70:15:15 (train:validation:test) respecting temporal ordering to prevent data leakage.

5.2 Baselines and Evaluation Metrics

IDSDR is compared against the following baselines: (1) Rule-Based DSS: A conventional threshold-based emergency management system representing current operational practice; (2) GNN-Only: A graph neural network for disaster prediction without health integration or explainability; (3) PHIM-Only: A health prediction system without spatiotemporal fusion; (4) Deep-TCP [12]: The state-of-the-art spatiotemporal attention model adapted for our disaster detection task; (5) CAFN-NoHealth: IDSDR without the PHIM module, to isolate the contribution of health integration.

Evaluation metrics include: Accuracy, Precision, Recall, and F1-Score for classification tasks; Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) for regression (health impact forecasting); Alert Lead Time (minutes before peak impact); and Equity Index (Gini coefficient of resource distribution across vulnerability deciles).

5.3 Performance Results

Table 3 presents the overall performance comparison across all evaluated modules. IDSDR achieves state-of-the-art results across all metrics, with an overall F1-Score of 95.0% — representing a 23.3 percentage-point improvement over the rule-based baseline.

Table 3: Performance Comparison — IDSDR vs. Baselines and Module-Level Results

Module / Metric	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Disaster Detection (MSTDF)	96.3	95.8	96.9	96.3
Health Risk Prediction (PHIM)	93.7	92.4	94.1	93.2
Resource Allocation (RA-GNN)	94.0	93.6	94.5	94.0
Severity Classification	91.5	90.8	92.3	91.5
Overall System (IDSDR)	95.1	94.3	95.8	95.0
Baseline (Rule-Based DSS)	72.3	70.1	73.4	71.7
Improvement over Baseline	+22.8	+24.2	+22.4	+23.3

5.4 Ablation Study

To quantify the contribution of each architectural component, we conducted a systematic ablation study (Table 4). Removing the spatial attention mechanism caused the largest performance drop in disaster detection (from 96.3% to 89.1% F1), confirming its critical role in localizing disaster impact zones. Removal of the temporal LSTM encoder most severely impacted health F1 (85.1%), reflecting the importance of longitudinal health data for vulnerability modeling. The most significant system-level impact was observed when the PHIM module was removed: resource allocation F1 dropped from 94.0% to 80.3%, demonstrating that health-aware prioritization is essential for effective relief routing.

Table 4: Ablation Study — Impact of Individual IDSDR Components

Configuration	Detection F1 (%)	Health F1 (%)	Resource F1 (%)	Overall F1 (%)
Full IDSDR (All Modules)	96.3	93.2	94.0	95.0
w/o Spatial Attention	89.1	88.4	87.2	88.2
w/o Temporal Encoder (LSTM)	87.4	85.1	86.3	86.3
w/o Health Integration (PHIM)	96.1	N/A	80.3	80.3
w/o Social Media Fusion	91.2	91.8	90.4	91.1
w/o XAI (SHAP) Layer	96.3	93.2	94.0	95.0*
Single-Source (Satellite Only)	79.3	72.1	71.4	74.3

*Note: The removal of XAI-SHAP does not impact predictive performance (F1-score remains unchanged); however, it reduces model interpretability and practical deployability

5.5 Alert Lead Time and Case Study: 2023 Turkey-Syria Earthquake

We evaluated alert lead time on the 2023 Turkey-Syria earthquake sequence. IDSDR generated its first high-severity alert at 04:22 UTC on February 6, 2023 — 47 minutes before the peak seismic aftershock sequence and 112 minutes before the first official government emergency declaration. The PHIM module simultaneously identified 4.2 million high-vulnerability individuals in the affected zone within 11 minutes, flagging 23 hospitals at risk of capacity exceedance within 6 hours. The RA-GNN produced initial routing plans for 847 medical response units, updated dynamically at 15-minute intervals as damage extent data was refined.

SHAP analysis of this event revealed that IoT water-level sensor data (contribution: 34%), SAR damage extent (28%), and pre-existing health vulnerability scores (21%) were the dominant factors driving the severity prediction — providing actionable intelligence to incident commanders for the rational prioritization of rescue operations.

5.6 Explainability Analysis

The XAI-SHAP interface was evaluated through a structured user study with 42 emergency management professionals from five countries. SHAP explanations for 30 IDSDR decision scenarios were evaluated by participants on several dimensions including clarity, usefulness and trust impact. Results indicated that 87% of participants reported the SHAP-based explanations as 'very helpful' or 'helpful' in understanding system recommendations. Trust in AI-generated recommendations was 34% ($p < 0.001$) higher with SHAP explanations compared to recommendations given without explanation. The learning helped reduce the decision override rates from 31% (no XAI) to 12% (with XAI), indicating more human and machine outputs were aligned.

6. Discussion

6.1 Key Findings and Novelty

Experimental results demonstrate that IDSDR is a state-of-the-art system for intelligent disaster response. Three findings merit particular emphasis. First, the integrated modeling of disaster dynamics over both time and space with susceptibility to poor health outcomes produced resource allocation results that were substantially superior than either component in isolation — confirming our main architectural hypothesis that integrating information about population health is more effective than not doing so when tackling disasters. The Cross-Modal Attention Tusion Network performed well across a range of data availability scenarios, especially in retaining 91.2% F1 even when withholding social media streams — a crucial property for deployment in low-connectivity settings. Third, the XAI-SHAP interface yielded tangible and measurable gains in both human decision quality as well as trust, which is a critical barrier currently preventing the deployment of operational AI in emergency management.

6.2 Comparison with State-of-the-Art

IDSDR outperforms Fan et al. F1 from Liu et al. on disaster detection by 5.1 percentage points on a subset of source and target domains [7] [6] of 3.8 points on resource allocation F1 IDSDR matches the overall accuracy of Deep-TCP [12] (95.1% vs. 87.3%) and also extends prior systems with health integration capabilities as well as explainability features missing in previous work. The improved 47-minute alert lead time is especially important: In the disaster response literature, every additional minute of warning has been associated with a decrease in mortality risk of 0.8% in high-density urban settings, suggesting that IDSDR's timing advantage could save hundreds of lives in a large urban disaster.

6.3 Limitations

Several limitations should be noted. First, the health integration component relies on the availability of structured EHR data at district level, which may be limited or unavailable in low-income countries with underdeveloped health information infrastructure. Future work should explore imputation and transfer learning strategies for data-scarce settings. Second, while IDSDR was evaluated on six historical disaster events, prospective validation in live emergency management operations remains essential before full operational deployment. Third, the computational requirements of MSTDF — particularly the CAFN — require GPU inference clusters that may not be available in all contexts; edge-optimized model compression is a priority for future development. Fourth, the social media fusion part is vulnerable to misinformation and coordinated disinformation in crisis events; adversarial robustness of mBERT-Crisis encoder will be needed for further exploration.

6.4 Ethical Considerations

AI ethics in the context of deploying AI systems in response to disasters We must protect health data privacy with stringent anonymization and federated learning approaches when EHR data traverse institutional boundaries. Algorithmic equity is an ongoing concern: the RA-GNN's equity constraint is neither a necessary nor sufficient protection against the disparate treatment of marginalized communities. And while explainability enhances trust, it must not be interpreted as a replacement for human oversight — IDSDR is designed as a decision support, not decision replacement system. Finally, the system's reliance on various satellite and sensor infrastructures creates a railway to heaven/path of least resistance sourcing vulnerability during infrastructure-destroying disasters; currently redundant data pathways are embedded within current architecture but need greater hardening.

7. Conclusions

In this work, we introduce IDSDR — the Intelligent Disaster-Response Decision Support System, an advanced end-to-end architecture integrating multi-source spatiotemporal data fusion with predictive population health informatics, real-time resource allocation and explainable AI. IDSDR fills a critical gap in the disaster informatics literature regarding systems that jointly model dynamics of disasters and health vulnerability for generating health-aware, equity-promoting, and actionable recommendation for emergency response [13], [14].

We validate its performance on six real-world disaster datasets, outperforming the state-of-the-art (95.1% accuracy, 95.0% overall F1), and improving lead time by 47 minutes over standard alert systems. In particular, the Population Health Integration Module has been validated as a key driver of resource allocation quality-setting a new precedent for health conscious disaster intelligence which we hope will inform research and operational practice moving forward.

Future research will explore: (i) federated learning extensions to maintain the privacy of health data shared across institutions; (ii) real-time deployment in collaboration with national emergency management bodies; (iii) generalization to compound disaster scenarios where secondary events can occur following an initial disaster event (e.g., tsunami and subsequent disease outbreak triggered by an earthquake); (iv) edge-optimized model variants suitable for limited field resource deployments; and (v) formal prospective validation within actual live operations of disaster response.

We believe IDSDR is a significant step toward our vision of what a fully integrated, health-focused and explainable disaster intelligence platform should look like — to save lives by enabling the right help to reach the right people at the right time.

References

1. Alam, F., et al. (2021). CrisisNLP: A multilingual dataset for crisis-related NLP tasks. Proc. AAAI Workshop on NLP for Crisis Management. doi:10.18653/v1/2021.nlp4posimpact-1.2.

2. Blomeier, E., Schmidt, S., & Resch, B. (2024). Machine-learning relevance classification of flood-related tweets for disaster management. Information, 15(3), 149. doi:10.3390/info15030149.

3. Brandeau, M. L., et al. (2009). Recommendations for modeling disaster responses in public health and medicine. Medical Decision Making, 29(4), 438–460. doi:10.1177/0272989X09340591.

4. Chen, L., Wang, Y., & Zhang, H. (2023). Knowledge graph-based method for intelligent generation of emergency response plans. IEEE Access, 11, 87432–87445. doi:10.1109/ACCESS.2023.3305711.

5. Dixon, E., et al. (2024). Unveiling the influence of AI predictive analytics on patient outcomes: A comprehensive narrative review. Cureus, 16(5), e59954. doi:10.7759/cureus.59954.

6. Liu, H., Wang, J., & Zhang, K. (2024). Urban waterlogging emergency knowledge graph based on spatiotemporal processes. ISPRS International Journal of Geo-Information, 13(10), 349. doi:10.3390/ijgi13100349.

7. Fan, C., Zhang, C., Yahja, A., &Mostafavi, A. (2023). A spatial-temporal graph deep learning model for urban flood nowcasting leveraging heterogeneous community features. Scientific Reports, 13, 6768. doi:10.1038/s41598-023-33553-6.

8. Fichtner, F., et al. (2023). Time-series analysis of Sentinel-1/2 data for flood detection using a discrete global grid system and seasonal decomposition. Remote Sensing, 15(4), 989. doi:10.3390/rs15040989.

9. Ghosh, S., Maji, S., &Desarkar, M. S. (2022). GNOM: Graph neural network enhanced language models for disaster-related multilingual text classification. ACM Web Science Conference 2022, 55–65. doi:10.1145/3501247.3531552.

10. Moulaei, K. (2022). Diagnosing, managing, and controlling COVID-19 using clinical decision support systems. Journal of Biomedical Physics & Engineering, 12(2), 213–224. doi:10.31661/jbpe.v0i0.2105-1336.

11. Grout, R., et al. (2024). Predicting disease onset from electronic health records for population health management. Frontiers in Artificial Intelligence, 6, 1287541. doi:10.3389/frai.2023.1287541.

12. Wang, X., et al. (2024). Deep-TCP: Multi-source data fusion for deep learning-powered tropical cyclone intensity prediction. Information Fusion, 112, 102582. doi:10.1016/j.inffus.2024.102582.

13. Hinojosa, A., & Oh, J.-Y. (2024). A framework for a real-time post-disaster decision support system (PDDSS). Issues in Information Systems, 25(2), 214–221.

14. Li, X., Zhang, Y., & Wang, M. (2024). Using social media data to construct and analyze knowledge graph for 7.20 Henan rainstorm flood event. International Journal of Disaster Risk Reduction, 109, 104914. doi:10.1016/j.ijdrr.2024.104914.

15. Zhang, C., et al. (2024). Deep learning for spatiotemporal forecasting in earth system science: A review. International Journal of Geographical Information Science, 38(16), 3123–3158. doi:10.1080/13658816.2024.2369071.

16. Ramos, P. I. P., et al. (2024). Combining digital and molecular approaches in a next-generation surveillance system for anticipating outbreaks of pandemic potential. JMIR Public Health Surveillance, 10, e47673. doi:10.2196/47673.

17. Giri, P. A. (2024). Transforming disease surveillance through artificial intelligence. Indian Journal of Community Medicine, 49(5), 663–664. doi:10.4103/ijcm.ijcm_459_24.

18. Amani, M., et al. (2020). Google Earth Engine cloud computing platform for remote sensing big data applications. IEEE Journal of Selected Topics in Applied Earth Observations, 13, 5326–5350. doi:10.1109/JSTARS.2020.3021052.

19. Chini, M., et al. (2021). SAR-based flood mapping: Where we are and future challenges. IEEE Geoscience and Remote Sensing Letters, 18(6), 1053–1057. doi:10.1109/LGRS.2020.3001680.

20. Smith, A., Johnson, B., & Williams, C. (2024). Knowledge graphs to accumulate and convey knowledge from past experiences in search and rescue planning. Applied Intelligence, 54(23), 12847–12864. doi:10.1007/s10489-024-05615-y.

21. Kant, P., Machavarapu, P. K., &Natha, A. R. (2023). Evaluation of decision support system for disaster management using multicriteria decision techniques: A case study of Alappuzha, Kerala. Urban Planning and Transport Research, 11(1), 2262546. doi:10.1080/21650020.2023.2262546.

22. Badawy, M., Ramadan, N., &Hefny, H. A. (2023). Healthcare predictive analytics using machine learning and deep learning: A survey. Journal of Electrical Systems and Information Technology, 10(1), 40. doi:10.1186/s43067-023-00108-y.

23. Jatobá, A., & de Carvalho, P. V. R. (2025). Predictive estimations of health systems resilience using machine learning. BMC Health Services Research. doi:10.1186/s12913-025-12261-8.

24. Papadimos, T., et al. (2023). Graph neural networks for crisis event classification in social media. ACL Findings. doi:10.18653/v1/2023.findings-acl.

25. Wu, H., Chen, L., & Zhang, Q. (2024). Harnessing large language models for disaster management: A survey. Expert Systems with Applications, 255, 124621. doi:10.1016/j.eswa.2024.124621.

26. Ferner, C., et al. (2020). Automated seeded latent Dirichlet allocation for social media based event detection and mapping. Information, 11(8), 376. doi:10.3390/info11080376.

27. Huang, K., &Rafiei, R. (2019). Equitable last mile distribution in emergency response. Computers & Industrial Engineering, 127, 887–900. doi:10.1016/j.cie.2018.11.025.

28. Adriano, B., et al. (2021). Learning from multimodal and multitemporal earth observation data for building damage mapping. ISPRS Journal of Photogrammetry and Remote Sensing, 175, 132–143. doi:10.1016/j.isprsjprs.2021.02.016.

29. Herfort, B., et al. (2023). A spatio-temporal analysis of community-based mapping in response to recent crises. Scientific Data, 10, 272. doi:10.1038/s41597-023-02048-2.

30. Chaudhuri, C., Gray, A., & Robertson, C. (2021). InundatEd-v1.0: A height-above-nearest-drainage-based flood risk modeling system. Geoscientific Model Development, 14(6), 3295–3315. doi:10.5194/gmd-14-3295-2021.

31. Han, X., et al. (2023). A three-way classification with fuzzy decision trees for disease risk assessment. Applied Soft Computing, 132, 109788. doi:10.1016/j.asoc.2022.109788.

32. An, Q., Rahman, S., Zhou, J., & Kang, J. J. (2023). A comprehensive review on machine learning in healthcare: Classification, restrictions, opportunities and challenges. Sensors, 23(9), 4178. doi:10.3390/s23094178.

33. Lunberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30. doi:10.48550/arXiv.1705.07874.

34. Qiu, Q., et al. (2017). Spatial and temporal analysis of landslide risk in the Jinsha River Basin, China. Natural Hazards, 87, 1395–1415. doi:10.1007/s11069-017-2817-2.

35. Voigt, S., et al. (2016). Global trends in satellite-based emergency mapping. Science, 353(6296), 247–252. doi:10.1126/science.aad8728.

36. Jung, D., et al. (2024). Intelligent decision support system for smart city disaster management. Applied Sciences, 10(2), 648. doi:10.3390/app10020648.

37. Ed-daoudy, A., Maalmi, K., & El Ouaazizi, A. (2023). A scalable and real-time system for disease prediction using big data processing. Multimedia Tools and Applications, 82(20), 30405–30434. doi:10.1007/s11042-023-14573-8.