Professional paper

Artificial Intelligence or Human Decision-making: A Study on Crisis Response Following an Earthquake

Milić, I.[1], Bilić, K.[2], Janeš, H.[3]

[1]University of Applied Sciences Velika Gorica,

Velika Gorica, igor.milic@vvg.hr
[2]Zagreb University of Applied Sciences,

Zagreb, krunoslav.bilic@tvz.hr

[3]University of Applied Sciences Velika Gorica,

Velika Gorica, hrvoje.janes@vvg.hr

Abstract: Decision-making in earthquake crises requires rapid coordination, informed analysis, and ethical responsibility. This study explores the comparative strengths of artificial intelligence (AI) and human judgment in such contexts through a simulated earthquake response scenario. The simulation was conducted as part of the 2022 Civil Protection Exercise Plan of the City of Velika Gorica, supported by Velika Gorica University of Applied Sciences and the Directorate of Civil Protection, with the goal of strengthening the city's disaster response readiness. The case examines how AI could contribute to operational efficiency and pattern recognition, while also recognizing the enduring importance of human expertise in ethical decision-making and adaptive leadership. Findings indicate that a hybrid model combining AI-driven tools with human oversight offers the most resilient framework for crisis response. The study underscores the need for continued development of integrated decision-support systems grounded in practical exercises and cross-sector collaboration.

Keywords: Earthquake, decision making, crises coordination, AI analysis

1. INTRODUCTION

Effective decision-making during disaster scenarios, especially earthquakes, remains a central challenge for crisis management authorities. These high-pressure situations demand rapid situational analysis, resource allocation, and communication across multiple agencies all under conditions of severe uncertainty and operational disruption (Boin et al., 2005). Traditionally, such decisions have relied heavily on human judgment, drawing from experience, intuition, and institutional memory. However, advances in artificial intelligence (AI) have introduced new possibilities for supporting or even augmenting human decision-making in real-time disaster response (Turoff et al., 2004; Comfort et al., 2004).

1.1 Background and significance

This paper is a conceptual and application-oriented analysis. It does not present or analyze empirical data from any exercise. It discusses where AI systems could support human decision-making in earthquake response, using the 2022 Velika Gorica Civil Protection exercise only as contextual illustration. AI technologies, particularly those utilizing real-time data analytics, geospatial mapping, and predictive modeling, have shown promise in improving response speed and efficiency. However, AI systems frequently struggle to comprehend complex social contexts and subtle human interactions, exercise ethical judgment, or adapt flexibly to unstructured or unforeseen developments. These qualities remain the domain of human crisis leaders, who must navigate not only logistical constraints but also emotional, political, and moral dimensions of disaster response (Özcan, 2021; Coombs, 2015).

1.2 Aim and scope

The aim is to map potential AI decision-support functions to key decision points in earthquake response and to explain how these functions can complement human judgment. This article offers conceptual mapping and does not collect, process, or compare primary data; the Velika Gorica field civil protection exercise serves solely as contextual grounding for illustrative decision points, not as an analyzable dataset.
Research questions:

1. What are the strengths and limitations of AI-based decision-support systems in simulated earthquake response?
2. In what ways does human expertise continue to provide value in managing disaster scenarios?
3. How can AI and human decision-making be effectively integrated to improve overall crisis response capacity?

2. CONCEPTUAL FRAMEWORK

Earthquake response places decision-makers under severe time pressure while information is incomplete, contested, or rapidly changing. Public leaders must coordinate diverse organizations, justify actions to citizens and media, and manage ethical trade-offs that affect vulnerable groups. These conditions amplify known bottlenecks: cognitive overload, hesitancy, centralization that slows operational throughput, and friction at the interfaces between agencies and volunteers. Empirical narratives from the 2011 Van earthquake in Turkey describe early coordination failures due to the absence of a shared incident command system, over-centralized approvals that delayed field action, and difficulties integrating spontaneous volunteers dynamics that illustrate how institutional structure and political accountability shape on-the-ground decision speed and quality (Özcan, 2021; Boin, et al., 2005; Coombs, 2015).

Public sector crisis models therefore conceptualize decision-making not as a single act but as a sequence of tasks sense-making, deciding, implementing, communicating meaning, and learning performed under uncertainty and organizational constraint. Even well-informed choices face the practical challenge of materializing across heterogeneous agencies that resist information sharing or have incompatible coordination styles. Non-decisions (delays, deferrals, evasions) are endemic and can be as consequential as explicit choices. These insights situate "human factors" at the core of crisis governance: leadership style, delegation practices, and communication discipline affect outcomes as much as technical capacity. They also clarify why any computational aid must be embedded in organizational routines that respect legal mandates, political legitimacy, and professional ethics (Boin et al., 2005; Coombs, 2015; Özcan, 2021).

2.1 AI-based decision-support possibilities

Against this backdrop, artificial intelligence can contribute in speed, scale, and consistency to several crisis tasks. Decision-support systems can fuse multi-source inputs—seismic feeds, remote sensing, crowdsourced reports, infrastructure layers into shared operational pictures and prioritized task lists. Continues work on emergency information systems emphasizes that coordination improves when information timeliness, relevance, and accessibility are raised for all actors, reducing effort and enabling earlier convergence on workable plans (Comfort, et al., 2004). In parallel, the dynamic emergency response management information system (DERMIS) design principles argue for dynamic, distributed, and user-configurable information flows that support collaboration across organizational boundaries rather than top-down command alone (Turoff, et al., 2004).

For logistics and resource allocation, relief-chain models using fault-tree analysis and Failure Mode, Effects and Criticality Analysis (FMECA) identify failure points from assessment through mobilization, transport, staging, and last-mile delivery. Such models make explicit the dependencies that often remain tacit in practice and thereby surface levers for optimization. When paired with modern data sources, they can guide priority setting for assets, routes, and staffing; detect bottlenecks early; and support "what-if" exploration of cascading effects. In large-scale events, this systematic perspective helps ensure that the visibility provided by dashboards translates into decisions that move supplies and services to the right places at the right times (Kumar & Havey, 2013).

AI can also triage information overload. Natural-language processing may summarize field reports, classify citizen requests, and flag anomalies; computer vision can accelerate building-damage assessments from aerial imagery; predictive models can estimate level of infrastructure damage or shelter demand under alternative assumptions. Yet these capabilities remain bounded by training data, feature engineering, and interface design. They require continuous validation, explicit uncertainty communication, and governance that constrain unintended consequences. Prior research cautions that computational efficiency without human interpretive capacity risks misalignment with local conditions and social priorities (Comfort et al., 2004; Turoff et al., 2004).

2.2 Integrating human and AI contributions

A pragmatic approach is a human-in-the-loop architecture in which AI systems propose humans dispose. Information systems supply common operating pictures, ranked options, and traceable rationales; human authorities arbitrate ethical dilemmas, negotiate inter-agency trade-offs, and own accountability for final choices. This design aligns with public-sector crisis models that stress fast but defensible decisions and with organizational lessons from prior earthquakes, which show the costs of over-centralization and the need for disciplined delegation and coordination. In practice, the integration challenge is less about model accuracy than about workflow: who sees what, when, and with what authority to act. Effective configurations clarify roles, encode escalation rules, and build audit trails that support learning after the event (Boin et al., 2005; Özcan, 2021; Turoff et al., 2004; Comfort et al., 2004).

For communication, AI tools may help maintain a single authoritative voice by filtering misinformation, harmonizing updates, and routing messages to appropriate spokespeople, but strategic messaging, empathy, and legitimacy remain human responsibilities. Literature on crisis communication underscores that unmanaged information flows can deepen crises, and that trust depends on known relationships and credible messengers. AI therefore supports but does not substitute for operational plans developed before the event and executed by trained officials during response (Coombs, 2015; Özcan, 2021).

2.3 Implications for local earthquake exercises

When mapped to the functional blocks of a local earthquake exercise, alerting and sense-making, assessment and routing, triage and surge management, logistics and staging, public information, and after-action learning AI contribute primarily to accelerating situational awareness, stabilizing coordination through shared data, and optimizing resource flows. Human decision-makers retain primacy where conflicts arise, where political license is required, and where improvisation must reconcile competing mandates or community expectations. The practical objective is a hybrid arrangement in which AI raises the floor of routine performance and humans navigate the ceiling of ambiguity, ethics, and legitimacy (Comfort et al., 2004; Boin et al., 2005; Kumar & Havey, 2013; Özcan, 2021).

The frameworks and models referenced here provide the theoretical basis for identifying potential AI functions applicable to local crisis exercises, rather than for evaluating actual performance. This conceptual stance foregrounds design choices such as data governance, role definitions, interface design, and transparent record-keeping that determine whether AI support reduces friction or introduces new failure modes. It also foresees next steps for empirical work: simulation exercises that measure coordination latency, logistics throughput, and communication coherence with and without AI assistance, under the oversight of public authorities responsible for ethical and legal compliance (Turoff et al., 2004; Comfort et al., 2004; Kumar & Havey, 2013).

3. METHODOLOGICAL CONTEXT

The 2022 Civil Protection Exercise of the City of Velika Gorica is used solely as an illustrative frame that anchors examples in realistic operations and organizational roles. This paper draws conceptually on the structure and objectives of the 2022 Velika Gorica Civil Protection Exercise to illustrate how existing AI-based models could align with real operational goals. No original data was recorded, processed, or analyzed. Exercise involved local emergency services and civil protection stakeholders with logistics support from national level (Civil Protection Directorate -Ministry of Interior.)

3.1 Exercise as illustrative frame

Publicly described features of the exercise provide a practical scaffold for discussing where AI systems might assist crisis managers. The scenario described a destructive early-morning earthquake affecting population, critical infrastructure, mobility, utilities, and public services across Velika Gorica, with activation of local coordination structures and a field command post. Within this operational picture, the paper treats the exercise elements as *contextual anchors* not as datasets to show how information flows, logistics, triage, and communication tasks could be supported by AI-enabled tools under the authority of human decision-makers.

3.2 Conceptual mapping logic

The mapping follows a simple logic tailored to practice:

- Identify decision points implied by the objectives. For example, issuing alerts requires situation recognition and channel selection; evacuation requires corridor selection and dynamic rerouting; triage requires prioritization under surge; logistics requires allocation under constraints; public communication requires message harmonization and rumor control. These points reflect standard emergency-management tasks that recur in municipal earthquake response.
- Associate each decision point with candidate AI functions described in prior models. Examples include sensor/image fusion for situational awareness, traffic and network flow algorithms for routing, predictive scoring for resource allocation, natural-language processing for report summarization, and dashboarding for cross-agency visibility. The paper references these models as design templates rather than as tools that were executed in the exercise.
- Specify human oversight and organizational preconditions. For each AI function, the mapping enumerates decision rights, escalation rules, interoperability expectations, uncertainty communication, and audit requirements that ensure legality, legitimacy, and ethical appropriateness remain with human authorities. This preserves accountability and aligns with the public-sector mandate of the civil protection system.

3.3 Boundaries and non-claims

The approach intentionally avoids methodological forms associated with empirical research. It does not propose sampling frames, instrumentation plans, reliability checks, or inferential tests.

It does not compare measured outcomes across "human-only" and "AI-assisted" conditions. Any contrasts discussed in later sections are illustrative propositions that show how specific AI functions might influence timeliness, coordination, or information quality if embedded in the workflows represented by the exercise objectives. Where the paper references the exercise, it is to make the conceptual mapping concrete e.g., by locating an AI dashboard within a field command-post workflow or by trying triage support to medical surge management not to infer performance effects.

3.4 Candidate AI model classes and input data by function

To remove ambiguity about which models operate on which data, the table below enumerates plausible model families and typical inputs for each functional area. These are exemplars, not prescriptive choices, and should be piloted under local governance and validation (Sidahmed, 2024).

3.4.1 Situational awareness and damage assessment

- Model classes. U-Net/FPN/PSPNet for semantic segmentation; Mask R-CNN for instance-level extraction of collapsed elements; lightweight detectors such as YOLOv8/Detectron2 for corridor blockage cues in UAV streams; two-stage pre/post change-detection and domain-adaptation pipelines for transfer across sensors and locales (Kızılay, 2024; Cheng, 2024; Zheng, 2024).
- Data inputs. Orthomosaics from UAVs, high-resolution satellite imagery, SAR for night/cloud cover, baseline cadastral/GIS layers, and dynamic road-network status feeds (OECD, 2025; Zheng, 2024).

3.4.2 Resource allocation and logistics optimization

- Model classes. Reinforcement learning for dynamic dispatch and staging; multi-period optimization with equity constraints; multi-agent RL for distributed coordination, embedded on top of fault-tree and FMECA structures that pinpoint leverage points along the relief chain (Yu, 2021; Ahmad, 2025; Kumar i Havey, 2013).
- Data inputs. Staging-point inventories, zone-level demand forecasts, transport-graph status, crew availability, shelter capacity, replenishment lead times, and disruption indicators from the local relief network (Kumar i Havey, 2013; Yu, 2021).

3.4.3. Public communication and coordination support

- Model classes. Fine-tuned BERT-family classifiers for incident relevance and priority; abstractive/extractive summarization for field logs and citizen reports; taxonomy-guided event detection and misinformation flags to standardize routing across agencies (Khallouli, 2025; He, 2025; Carvalho, 2025).
- Data inputs. Structured incident logs, transcribed radio/field notes, geotagged social-media streams, and operator annotations for supervised updates (Carvalho, 2025; He, 2025).

3.4.4 Governance, transparency, and learning

- Design choices. Role-based access, mandatory decision logs, explicit uncertainty displays, provenance tracking, and data-protection controls aligned with public-sector audit requirements; DSS patterns that enable explainability and traceability across agencies (Comfort, et al., 2004; Turoff, et al., 2004; OECD, 2025).

4. APPLICATION AREAS OF AI IN EARTHQUAKE RESPONSE

4.1 Situational awareness and damage assessment

AI can accelerate early situational awareness by fusing heterogeneous data streams into a common operating picture. Computer-vision models applied to drone or satellite imagery can flag probable structural damage, blocked corridors, and emerging fire risks. Combined with sensor inputs and basic geospatial layers, these models can generate rapid heatmaps of impact and accessibility that guide initial tasking and alerting. Prior work on information flows in rapidly evolving disasters shows that timeliness, relevance, and shared access to curated information improve coordination and reduce duplicative effort, which are central goals for any AI-enabled dashboard.

4.2 Resource allocation and logistics optimization

Once the operating picture stabilizes, AI can support logistics by identifying high-leverage interventions across the relief chain. Fault-tree and Failure Mode, Effects and Criticality Analysis (FMECA) frameworks make interdependence explicit from assessment through mobilization, transport, staging, and last-mile delivery. Coupled with current demand estimates and infrastructure status, this perspective enables algorithmic prioritization of supplies, routing around degraded networks, and proactive detection of bottlenecks that often go unnoticed until they cascade. In large events, optimization informed by these models can reduce stockouts at critical nodes, align staff and assets with likely surge, and support "what-if" planning for aftershocks or secondary hazards.

4.3 Public communication and coordination support

Earthquake response routinely suffers when coordination is ad hoc, approvals are over-centralized, and volunteers operate outside a common system. Narratives from prior crises document how the absence of an incident command structure, combined with upward transfer of routine decisions, produced delays and public frustration. These dynamics imply clear opportunities for AI tools that triage and route information rather than replace authority: natural-language processing can summarize field reports, rank incoming citizen requests, and surface anomalies; cross-agency dashboards can expose task ownership and status to reduce duplication; and alerting assistants can harmonize messages for designated spokespersons to maintain a single authoritative voice. The target is disciplined coordination with humans in charge, not message automation at the expense of legitimacy.

4.4 Ethical and leadership considerations

AI systems should augment, not supplant, public leadership in ethically charged decisions, such as triage, evacuation prioritization, and risk communication. Crisis governance models emphasize sense-making under uncertainty, rapid yet accountable decisions, and coordinated implementation across heterogeneous organizations. In this setting, AI contributes auditable recommendations with uncertainty cues, while human authorities adjudicate trade-offs, ensure procedural fairness, and retain accountability. Over-centralization that slows field activities, or delegation without oversight, can both degrade outcomes; effective designs therefore pair AI transparency with clear decision rights, escalation rules, and post-event learning mechanisms.

Table 1: Human and AI-supported decision-making

Dimension	Human Decision-Making	AI-Supported Decision-Making
Timeliness	Moderate to delayed response due to communication constraints and manual assessments.	High-speed data processing enables near-instantaneous situational awareness (Comfort et al., 2004).
Accuracy	Based on expertise and judgment, occasional errors and overlaps in task allocation.	High consistency in analysis and task distribution; minimizes redundancy (Kumar & Havey, 2013).
Adaptability	Strong improvisational skills; effective in unpredictable human behavior (Boin et al., 2005).	Limited to training data; less capable of responding to unexpected human dynamics.
Coordination	Dependent on interpersonal networks; sometimes siloed and inconsistent.	Centralized dashboards ensure inter-agency synchronization and transparency (Turoff et al., 2004).
Ethical/Social Judgment	High sensitivity to ethical priorities and social needs (e.g., vulnerable groups) (Özcan, 2021).	Lacks emotional intelligence and ethical discernment cannot contextualize moral choices.

This functional reframing positions AI as a catalyst for speed, scale, and consistency in information processing and logistics, while reserving normative judgment, political legitimacy, and stakeholder engagement for human leaders. The categories above align with the operational tasks commonly activated in municipal earthquake exercises and can be used to scope pilot integrations that preserve accountability while testing where AI support yields measurable gains in coordination latency, throughput, and message coherence.

5. CONCLUSION AND RECOMMENDATIONS

This paper proposes that integrating AI decision-support functions with human oversight may strengthen earthquake crisis response. The proposition is conceptual: AI is mapped to decision points where speed, scale, and consistency in information processing and logistics matter, while ethical judgment, legitimacy, and accountability remain with human authorities. AI can contribute via data fusion, pattern recognition, and relief-chain optimization; humans arbitrate trade-offs, coordinate heterogeneous agencies, and lead public communication under uncertainty.

To operationalize this framework, implement narrowly scoped AI functions within controlled exercises, for example, imagery-based damage triage or routing optimizers under clearly defined decision rights; then evaluate their contribution by systematically tracking coordination latency, corridor clearance times, stockout rates at staging points, and cross-agency message coherence before considering wider deployment. Establish governance by design with human-in-the-loop protocols, model-output traceability, and explicit uncertainty cues, ensuring tools reinforce incident command and designated spokesperson roles rather than supplanting them. Build organizational readiness through targeted training on delegation, information-sharing agreements, and structured integration of spontaneous volunteers to avoid delays from over-centralization and ad hoc coordination. Prioritize interoperability and data stewardship policies that enable responsible access, retention, and interagency exchange even under degraded communications. Implement rigorous after-action learning by auditing AI-supported workflows post-exercise to refine protocols and model scopes without eroding legitimacy or ethical standards.

These steps provide a practical path to empirically test the framework and determine where AI augmentation yields measurable gains while preserving human control over consequential decisions.

6. REFERENCES

Ahmad, M. (2025). An enhanced deep reinforcement learning approach for humanitarian logistics. *Engineering Applications of Artificial Intelligence*.

Boin, A., ’t Hart, P., Stern, E., & Sundelius, B. (2005). *The politics of crisis management: Public leadership under pressure*. Cambridge, UK: Cambridge University Press.

Carvalho, P. (2025). A collaborative taxonomy of social media indicators for disasters. *PLOS ONE*.

Cheng, M. Y. (2024). Computer vision-based post-earthquake inspections for buildings. *Automation in Construction*.

Comfort, L., Ko, K., & Zagorecki, A. (2004). Coordination in rapidly evolving disaster response systems: The role of information. *American Behavioral Scientist*, 48, 295–313. https://doi.org/10.1177/0002764204268987

Coombs, W. T. (2015). *Ongoing crisis communication: Planning, managing, and responding* (4th ed.). Thousand Oaks, CA: Sage.

He, C. (2025). Social media analytics for disaster response. *Applied Sciences*.

Khallouli, W. (2025). Integrated deep learning for emergency tweet identification (BERT + rules). *Social Network Analysis and Mining*.

Kızılay, F. (2024). Evaluating fine-tuned deep learning models for real-time building damage detection from UAV imagery. *Discover Artificial Intelligence*.

Kumar, S., & Havey, T. (2013). Before and after disaster strikes: A relief supply chain decision support framework. *International Journal of Production Economics*, 145(2), 613–629. https://doi.org/10.1016/j.ijpe.2013.05.016

Organisation for Economic Co-operation and Development (OECD). (2025). *Using AI to measure disaster damage costs: Methodology and the example of the 2018 Sulawesi earthquake*. Paris, France: OECD Publishing.

Özcan, Ş. (2021). *Effectiveness in crisis management*. Ankara, Turkey: Iksad Publishing House.

Sidahmed, E. (2024). Strategic support systems for crisis management: A literature review. *Foundations and Trends in Information Systems*.

Turoff, M., Chumer, M., Van de Walle, B., & Yao, X. (2004). The design of a dynamic emergency response management information system (DERMIS). *Journal of Information Technology Theory and Application*, 5(4), 1–35.

Youngblood, S. (2010). Ongoing crisis communication: Planning, managing, and responding, 2nd edition (Coombs, W. T.) and *Handbook of Risk and Crisis Communication* (Heath, R. L., & O'Hair, H. D., Eds.) [Book reviews]. *IEEE Transactions on Professional Communication*, 53, 174–178. https://doi.org/10.1109/TPC.2010.2046099

Yu, L. (2021). Reinforcement learning approach for resource allocation in humanitarian logistics. *Expert Systems with Applications*.

Zheng, Z. (2024). Towards transferable building damage assessment via domain-adaptive change detection. *Remote Sensing of Environment*.