A Modular Simulation and Decision Platform for Enhancing Mountain Rescue Operations

Radoslaw Klimek
AGH University of Krakow
Krakow, Poland *rklimek@agh.edu.pl*

Krzysztof Łazarz
AGH University of Krakow
Krakow, Poland *kjlazarz@student.agh.edu.pl*

Michał Pałucki
AGH University of Krakow
Krakow, Poland *palucki@student.agh.edu.pl*

Mateusz Świątek
AGH University of Krakow
Krakow, Poland *mateoswiatek@student.agh.edu.pl*

Abstract

Mountain rescue operations face uncertainty, difficult terrain, and limited communication. This paper introduces a modular platform combining agent-based simulation in Unity with a microservice decision engine for real-time risk assessment. The system models weather, avalanches, tourist and animal movement, and processes sensor data using logic-based inference and machine learning. Early results confirm sub-second latency and high accuracy in hazard detection. Beyond technical capabilities, the platform proposes a cloud-based, data-driven business model designed for scalable deployment in rescue services, insurance, and tourism management.

1. INTRODUCTION

Mountain rescue operations are characterised by uncertainty, limited visibility, and rapidly changing environmental conditions. A harsh terrain and incomplete localisation of individuals in need frequently hamper rescue coordination and delay response. Traditional reactive approaches offer limited support in such dynamic settings. Advances in smart sensing, simulation technologies, and AI-supported reasoning have opened new possibilities for improving operational awareness and preparedness. However, most existing solutions focus on isolated components—such as geospatial modelling or weather forecasting—without combining real-time sensing, behaviour modelling, and decision logic in a unified system.

In this paper, we present a modular platform that addresses this gap through the combination of simulation and real-time decision support. The system couples a Unity-based environment–modelling weather, avalanche risk, and behaviour–linked to a microservice decision engine that transforms heterogeneous sensor data into operational recommendations. By combining logic-based inference with data-driven risk assessment, the platform supports proactive alerting and situational awareness. Beyond its technical architecture, the platform contributes a data-oriented business model that enables scalable deployment in mountain rescue, insurance, and tourism services. It is designed for both institutional use and integration with public-facing mobile infrastructure. Early tests confirm the system's responsiveness, modularity, and ability to reflect critical risk dynamics, supporting its role in future mountain SAR systems.

2. RELATED WORKS

The current platform builds on work [3] by transitioning from a unified simulation system to a modular architecture that separates simulation and decision making. The adoption of containerised microservices and AI-driven reasoning improves scalability and responsiveness in high-risk real-time environments. Nasar et al. [7] reviewed decision support systems in SAR, emphasising the growing use of GIS, remote sensing, and AI. They identified a gap in platforms combining live sensor data, agent-based modelling, and dynamic reasoning—an area our system directly addresses. Damaševičius et al. [1] examine the use of the Internet of Emergency Services (IoES) in disaster management, emphasizing real-time sensing and communication for faster and more coordinated responses.

Nakamachi et al. [6] introduced a multi-layered cellular automaton (CA) to simulate neurogenesis, demonstrating how layered CA structures can model complex spatiotemporal phenomena. This inspired our approach to weather simulation, where separate layers represent distinct meteorological variables such as temperature, wind, or humidity. Papić et al. [8] propose a GIS-based system using a Person Mobility Algorithm to define search zones in a mountainous terrain. While their method enhances SAR planning with real-time data, our platform advances further through agent-based simulation, sensor fusion, and AI-driven decision support within a modular architecture. Dragoni et al. [2] present a foundational overview of microservice architecture, highlighting its advantages over monolithic systems and identifying key challenges such as orchestration and service coordination, issues also central to our platform's design.

3. METHODOLOGY

The system is organised as a modular platform comprising two principal components: a Simulation Environment and a Decision-Making System. This dual-layered design supports both controlled experimentation and real-time operational response. Simulations generate training data, while decision logic processes live inputs into alerts. The architecture ensures scalability, interoperability with external systems (e.g., GIS, mobile infrastructure), and continuous model refinement. The platform was built using Java-based microservices with Docker, Postgres and RabbitMQ, integrated into a CI/CD development pipeline.

Decision Making System. At the core of the platform, the Decision-Making System ingests heterogeneous data streams from simulated and external sources—such as GPS, BTS, and animal telemetry—and converts them into actionable outputs. It is implemented as a set of containerised microservices, each addressing a specific stage of a data-processing pipeline.

The pipeline proceeds from raw data ingestion and validation, through transformation and tabular structuring, to reasoning and decision output. The reasoning engine employs both logic-based inference and machine learning to estimate risk levels and guide prioritised responses [5]. Figure 1 outlines this modular data flow.

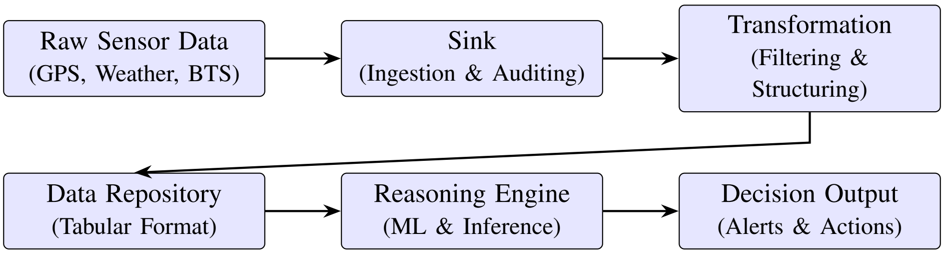

Fig. 1. Compact data pipeline of the Decision-Making System. Input data is processed through modular stages toward AI-supported recommendations.

Data Collection and Integration. A major challenge lies in aggregating and harmonising diverse data sources, particularly in sensor-driven environments where context-awareness and proactive processing play a key role [4]. The Sink module addresses this by validating, auditing, and standardising inputs from BTS triangulation, GPS devices, weather stations, and wildlife trackers.

Inputs are transformed into tabular formats with coordinates, vectors, metrics, and timestamps. Standardised JSON schemas ensure seamless microservice communication. Once stored, the data is accessible to downstream reasoning processes for predictive analysis and warning generation.

Reasoning. The reasoning component interprets structured data to assess risks such as distress incidents or environmental hazards. Unlike rigid, rule-based methods, the system utilises a dynamic, data-driven architecture that adapts to real-world variability.

Key stages include feature normalisation for input standardisation, neural inference trained on historical and simulated data, and probabilistic alerting based on threshold assessments. This approach reduces manual configuration and improves responsiveness. Future work will explore transfer learning and uncertainty estimation to enhance generalisability and robustness.

Simulation Environment. Implemented in Unity, the simulation environment supports real-time, agent-based modelling. Its modular design enables the representation of tourists, animals, weather, and avalanche conditions. Unity's scripting and visualisation capabilities also facilitate scenario configuration and interactive testing. Agent-based modelling is used for representing tourists and animals as autonomous entities, enabling realistic movement and interaction. In parallel, weather dynamics are approximated using a multi-layered cellular automaton framework.

Multi-Layered Weather Cellular Automata. A multi-layered cellular automata (CA) model is used to simulate weather phenomena such as temperature, wind, humidity, and rainfall. Each meteorological factor is modelled in a separate layer over a spatial grid. Although conceptually clear and modular, CA's computational cost suggests future hybridisation with empirical data and stochastic models.

Tourist and Animal Models. Tourists are represented as autonomous agents navigating OpenStreetMap-based routes, influenced by traits such as group size and behavioural tendencies. Animal agents—primarily bears—are simulated using GPS data and proximity rules. The models provide rich interaction data for downstream reasoning.

Avalanche Model. The avalanche module integrates terrain analysis, snowpack modelling, and triggering conditions to estimate hazard levels. Influenced by platforms, for instance Skitourenguru, the model includes terrain analysis (e.g., slope and aspect), meteorological inputs that drive snow dynamics, trigger thresholds for risk evaluation, and agent patrols responsible for reporting hazardous zones.

The model allows UI-based configuration and future hybrid integration. Figure 2 illustrates the model's workflow.

4. EVALUATION

The platform demonstrates the viability of integrating simulation and AI-driven decision-making for mountain rescue. Core components are functional, supporting real-time data ingestion, processing, and visualisation, and providing a basis for future predictive capabilities.

Simulator Results. Initial simulations used a bounded rectangular map, which proved limiting. The system was extended to a global context using WGS-84 geospatial referencing. A multi-layered cellular automata (CA) model was implemented, with early visualisation tools aiding validation.

The tests showed that CA is insufficient for capturing complex weather dynamics. While decision-making could rely solely on weather station data, simulated weather may enhance spatial and temporal resolution for incident prediction.

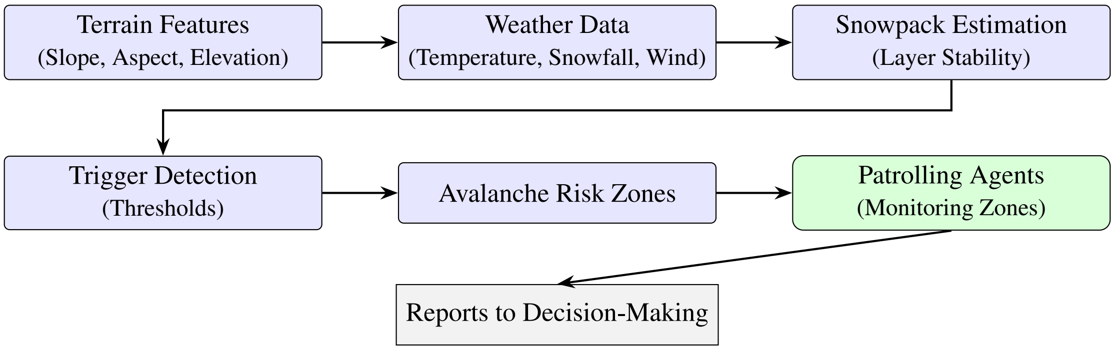

Fig. 2. Avalanche Model workflow: Terrain and weather data inform snowpack estimation and triggering conditions; resulting risk zones are monitored by patrolling agents who issue reports to the decision-making system.

Simulated sensor data achieved sub-second ingestion times (for <2kB packets). In synthetic avalanche scenarios, the reasoning module produced early warnings with 95% accuracy relative to predefined thresholds.

Decision Making System Results. The system uses containerised services orchestrated via Docker for maintainability and portability. Key infrastructure includes: Caddy (web proxy), HAProxy (load balancing), RabbitMQ (messaging), Postgres and MongoDB (data storage), Loki and Grafana (monitoring and logs).

Domain logic is delivered via Maven-based Java services. Three docker-compose configurations (development, staging, production) support modular deployment. Despite initial setup overhead, containerisation has streamlined integration and multi-environment deployment.

5. BUSINESS MODEL AND VALUE CREATION

The proposed platform establishes a hybrid business model that integrates licensed deployment for institutional users with the possibility of selective public access. Its cloud-native architecture allows scalable, multi-region operation, enabling multiple rescue organisations to monitor distinct areas simultaneously. Preliminary deployment scenarios indicate low infrastructure cost, with optional extensions through mobile or cloud services, subject to stakeholder-specific needs.

Licensing and Deployment Scenarios. Rescue agencies can access the system via license agreements, enabling configuration and monitoring of defined operational zones with full simulation, sensor integration, and decision support. A parallel public-facing version—with reduced resolution or limited functions—can serve as a safety-oriented public good. A key challenge is balancing data granularity and accessibility across user tiers.

Integration with Insurance Services. The system is well-positioned to support partnerships with the insurance sector. In many mountainous regions, accident insurance is legally mandated for tourists. By consenting to real-time monitoring within designated zones, individual users could benefit from reduced premiums or expanded coverage during their stay. This model aligns with usage-based insurance paradigms, in which behavioural transparency and environmental context inform about dynamic pricing.

Collaboration with the Tourism Sector. The platform produces behavioural data on tourist movement, group dynamics, route choices, and hazard exposure. These datasets assist tourism authorities in infrastructure planning (e.g., trail maintenance), targeted risk communication, and forecasting visitor flows and threats. Applying machine learning to historical and real-time data reveals latent patterns, enabling evidence-based interventions and improved visitor safety and experience.

Cloud-Based Operational Model. The system is designed for elastic deployment in cloud environments, enabling updates, scalability, and high availability, supporting efficient service delivery across multiple mountain regions and organisations, minimising on-premise maintenance while maintaining data sovereignty and access control for each client.

Value Proposition Summary. By transforming raw environmental and behavioural data into actionable insights, the platform benefits multiple stakeholders. Rescue services gain early threat detection and optimised deployment; insurers access real-time risk profiles and behaviour-linked pricing; tourism operators receive tools for safety and planning; and tourists enjoy improved safety and guidance. The platform illustrates how smart environments and AI can reshape business models in high-risk natural contexts, shifting from reactive response to proactive, data-informed ecosystem management.

6. CONCLUSIONS

This work introduced a modular platform combining simulation, real-time data processing, and automated reasoning to support mountain rescue. The system enables risk modelling and alert generation through the scalable integration of environmental and behavioural data. The preliminary tests confirm low-latency processing and modular deployment. Future work will enhance predictive models and validate the platform in operational settings.

References

[1] Damaševičius, R., Bacanin, N., and Misra, S.: From Sensors to Safety: Internet of Emergency Services (IoES) for Emergency Response and Disaster Management. In: *Journal of Sensor and Actuator Networks* 12.3 (2023), p. 41.

[2] Dragoni, N., Dustdar, S., Larsen, S., Mazzara, M., Montesi, F., Mustafin, R., and Safina, L.: "Microservices: Yesterday, Today, and Tomorrow". In: *Present and Ulterior Software Engineering*. Ed. by Broy, M., Ernst, E., and Wirsing, M. Springer, 2017, pp. 195–216.

[3] Klimek, R.: Ambient-aware continuous aid for mountain rescue activities. In: *Information Sciences* 653 (2024), 31 pp. URL: https://www.sciencedirect.com/science/article/pii/S0020025523013579.

[4] Klimek, R.: Sensor-enabled context-aware and pro-active queue management systems in intelligent environments. In: *Sensors* 20.20 (2020), pp. 1–29.

[5] Klimek, R. and Kotulski, L.: Proposal of a multiagent-based smart environment for the IoT. In: *Workshop Proceedings of the 10th International Conference on Intelligent Environments, Shanghai, China, 30th June–1st of July 2014*. Ed. by Augusto, J. C. and Zhang, T. Vol. 18. Ambient Intelligence and Smart Environments. IOS Press, 2014, pp. 37–44.

[6] Nakamachi, E., Nakayama, A., Yamamoto, T., Morita, Y., and Sakamoto, H.: Development of a Novel Simulation Code to Predict Three-Dimensional Neurogenesis by Using Multilayered Cellular Automaton. In: *International Journal of Computational Methods and Experimental Measurements* 7 (2019), pp. 201–211.

[7] Nasar, W., Nazir, M., Habib, H., and Kim, K.-H.: The Use of Decision Support in Search and Rescue: A Systematic Literature Review. In: *ISPRS International Journal of Geo-Information* 12.5 (2023), p. 182. URL: https://www.mdpi.com/2220-9964/12/5/182.

[8] Papić, V., Gudelj, A. Š., Milan, A., and Miličević, M.: Person Mobility Algorithm and Geographic Information System for Search and Rescue Missions Planning. In: *Remote Sensing* 16.4 (2024), p. 670. URL: https://www.mdpi.com/2072-4292/16/4/670.