

AURORA: AI-supported Unified Risk Observation and Response Assistant

, CLEMENS HÖLLER, , and , WU Vienna
University of Economics and Business, Austria

1 Concept and Concrete Objectives

Problem Context. The increasing frequency and severity of extreme weather events expose significant limitations in current emergency alert systems such as AT-Alert. Climate change is intensifying extreme weather phenomena and increasing the need for reliable early warning systems [15, 35]. These systems are primarily designed for rapid, large-scale dissemination of warnings, relying on broadcast-based approaches that deliver identical messages to all users within a geographic region [13, 31]. While effective in speed and coverage, this approach leads to major shortcomings, including over-alerting, lack of contextual relevance, and reduced user trust [25, 37]. As a result, users often receive warnings that are not directly relevant to their situation, contributing to alert fatigue and decreasing responsiveness in critical situations [7, 39].

Current research efforts address individual aspects of this problem, such as improving location-based targeting or applying machine learning for classification and prioritization [22, 32]. However, these solutions remain fragmented and fail to provide a unified system that integrates data, risk evaluation, and communication [8, 37]. In particular, existing systems lack the ability to combine structured knowledge, geospatial reasoning, and real-time data into a coherent framework that enables personalized and context-aware emergency alerts.

AURORA System Concept. To address these limitations, this project proposes AURORA (AI-supported Unified Risk Observation and Response Assistant), an integrated system that enhances existing emergency alert infrastructures through a three-layer architecture. The core idea is to transform emergency alerting from a generic broadcast model into a context-aware, risk-based, and personalized communication system.

The first component of the system is a Knowledge Graph layer, which integrates heterogeneous data sources such as geospatial information, infrastructure data, and real-time hazard events. Knowledge graphs have been identified as a powerful approach for structuring and integrating heterogeneous disaster-related data [34, 42]. The knowledge graph provides a structured and interconnected representation of the environment, capturing not only entities such as locations, hazards, and infrastructure, but also their relationships and dependencies. This enables the system to reason about how different elements interact and what is relevant in a specific spatial and temporal context.

Building on this, the Risk Assessment layer evaluates the actual impact of hazards on specific users. By combining geospatial analysis with real-time data, the system determines whether a user is affected by a given event and classifies the level of risk (e.g., low, medium, high). Geospatial risk modeling approaches have proven effective for identifying dynamic hazard zones and supporting decision-making in crisis scenarios [23, 33]. This layer introduces an explicit and explainable decision-making step, ensuring that alerts are based on transparent criteria rather than implicit assumptions.

Finally, the Grounded LLM layer translates structured risk information into natural language alerts. Recent advances in large language models demonstrate strong potential for generating context-aware and decision-supportive communication, particularly when grounded in structured data sources [4, 9]. Unlike traditional approaches, the large language model operates on verified and structured input, ensuring that all generated messages are grounded in factual data.

Authors' Contact Information: ; Clemens Höller; ; , WU Vienna University of Economics and Business, Vienna, Austria.

This reduces hallucination risks and enables the generation of clear, context-aware, and actionable warnings tailored to the user's situation [5].

Contribution and Objective. The key contribution of this concept lies in the integration of structured knowledge representation, geospatial reasoning, and controlled language generation into a single system. By combining these components, the project moves beyond isolated improvements and introduces a unified approach that enables precise, explainable, and personalized emergency communication. Ultimately, this approach aims to improve the effectiveness of public warning systems, reduce unnecessary alerts, and strengthen user trust and response in critical situations [31, 37].

Concrete Objectives. To achieve this goal, the project defines the following concrete objectives:

Objective 1: Build a queryable emergency knowledge graph

Develop, by month 7, an RDF-based emergency knowledge graph integrating at least three heterogeneous data sources geospatial, real-time hazard, and infrastructure data and enabling SPARQL-based identification of hazard-location-infrastructure relationships. Success will be measured against a predefined catalogue of at least 20 competency questions, with the target that at least 95% of these questions are answered correctly.

Objective 2: Implement dynamic user-specific risk assessment

Implement, by month 10, a spatio-temporal risk assessment module that classifies users into low, medium, and high risk levels based on their location, hazard exposure, infrastructure context, and temporal conditions. The module will update risk zones within 2 minutes of new hazard data. Evaluation will be conducted on simulated users within defined hazard scenarios, assessed against an independent expert-annotated ground truth.

Objective 3: Validate personalized alerts against broadcast-based baselines

By month 18, integrate all components into a working prototype and evaluate it in at least three realistic emergency scenarios. The system will be compared against a reconstructed region-based broadcast baseline modeled on AT-Alert, using defined metrics including precision and recall of affected-user identification, the false-alert rate, and the proportion of irrelevant alerts per scenario.

2 State of the Art

Limitations of Current Emergency Alert Systems. Current Emergency Alert Systems such as AT-Alert in Austria are designed around a single priority: getting a warning out quickly to as many people as possible. Broadcast-based infrastructures make this feasible, but they also make personalization nearly impossible [31, 37]. Messages are identical for all recipients, regardless of where they are, how close they are to the actual danger or whether they are directly affected at all. This creates a fundamental mismatch between how alerts are delivered and how risk is experienced in real-world situations [25, 30].

This mismatch reflects a central limitation: the lack of precise location awareness and contextual understanding. Existing systems typically operate on coarse geographic regions and cannot determine a person's exact position relative to a dynamic threat [11, 29]. At the same time, risk is not modeled as a gradient. There are no differentiated risk zones that distinguish between high, medium or low exposure within the same event. As a result, systems cannot adapt either the urgency or the content of alerts [23, 33]. Individuals inside a high-risk area may receive the same message as those on the periphery or completely outside the danger zone.

These limitations lead to well-documented consequences. Over-alerting is one of the most persistent problems, where users receive frequent notifications that are not relevant to their situation [7, 39]. This contributes directly to alert fatigue, where repeated exposure to low-value alerts causes users to become desensitized and eventually ignore warnings altogether [31, 39]. Trust in the system declines, and response behavior weakens [25, 30]. In addition, most

alerts remain generic and lack actionable guidance. They provide high-level warnings but do not translate risk into concrete, situation-specific instructions, leaving users uncertain about how to respond.

Isolated Advances in Existing Approaches. Research has addressed these shortcomings by improving individual components of emergency alert systems. Location-aware approaches attempt to improve geographic relevance using geofencing and spatial filtering [10, 32]. While this allows alerts to be restricted to certain areas, these approaches remain static and do not account for dynamic changes in risk. Machine learning methods have been introduced for event detection and prioritization, using classification models to process data from sensors and social media in near real time [22, 24]. Natural language processing techniques extend this further by extracting crisis-relevant information from unstructured sources such as social media posts and incident reports [6, 41]. In parallel, advances in data integration enable the combination of heterogeneous data sources, including meteorological data, infrastructure data, and satellite imagery, to improve situational awareness.

While these developments are valuable, they remain largely disconnected. Existing systems typically focus on either data collection, analysis or communication, but do not integrate these components into a unified architecture [30, 34]. As a result, even though the necessary capabilities exist, they are not combined in a way that enables context-aware and personalized alerting [37].

Knowledge Graphs and Risk Assessment. Knowledge Graphs provide a promising foundation for addressing this fragmentation by acting as a structured and centralized representation of heterogeneous data. In disaster management, knowledge graphs model entities such as hazards, locations, infrastructure, populations and resources, as well as the relationships between them [34, 42]. These relationships capture spatial, temporal and causal dependencies, enabling a more comprehensive understanding of complex situations [12, 38]. Most implementations rely on Semantic Web technologies, particularly the Resource Description Framework (RDF) and ontology-based modeling, which allow data to be represented as triples and linked across domains [21, 26]. This enables interoperability, flexible integration and advanced querying through languages such as SPARQL.

Several projects illustrate the potential of this approach. The XR4DRAMA project uses an RDF-based knowledge graph to integrate diverse disaster-related data into a unified semantic structure, focusing on interoperability and data consistency [40]. The CRISP case study applies similar principles in real-world scenarios, integrating incident reports and operational data to support querying and analysis during crisis response [3]. More recent work explores multimodal knowledge graphs that combine textual, sensor and geospatial data, as well as graph-based analytical methods for identifying patterns and relationships [17, 38]. These approaches demonstrate that knowledge graphs can serve as a single source of truth, bringing together fragmented data into a coherent and structured representation.

However, knowledge graphs alone do not address the challenge of translating data into actionable communication. Most implementations remain focused on analysis and situational awareness rather than user-facing outputs [34, 42]. They provide a strong foundation for understanding a crisis but do not bridge the gap to personalized alerts.

To move from data to action, research has introduced risk assessment approaches that derive risk levels from structured data. These methods use rule-based logic and analytical models to evaluate factors such as hazard intensity, proximity and environmental conditions [18, 33]. Concretely, these approaches encode expert knowledge as explicit decision rules: spatial operations such as buffer and proximity analysis and spatial overlays are combined with threshold values (e.g. assigning a high risk level when hazard intensity exceeds a defined limit within a given distance), commonly implemented in GIS frameworks such as QGIS or PostGIS [2]. This enables the creation of dynamic risk zones, where different areas are assigned different levels of severity. Such segmentation is a key prerequisite for personalization, as it

allows users to be grouped based on their exposure to risk. However, these approaches are often limited in scope, lack integration with broader systems and do not extend into communication layers [23].

LLMs and Emerging Integrated Approaches. Large Language Models introduce new possibilities for addressing the communication side of this problem. Models capable of processing and generating natural language at scale are increasingly used in decision support contexts, including crisis management [4, 20]. Retrieval-Augmented Generation improves these capabilities by combining generative models with external knowledge retrieval, enhancing factual grounding and contextual relevance [19, 28]. In emergency contexts, LLMs have been used to process unstructured data such as social media, reports and communication logs to support situational awareness [22, 27].

At the same time, the limitations of LLMs in safety-critical environments are well established. LLMs generate outputs based on statistical patterns rather than true understanding, which can lead to hallucinations and unreliable information [5]. This is particularly problematic in emergency scenarios, where incorrect outputs can have serious consequences. Furthermore, LLMs are typically optimized for generating descriptive text rather than concise, actionable instructions. They also lack native integration with structured data, real-time inputs and geospatial reasoning, which limits their applicability in dynamic alerting systems [4].

Recent research has begun to combine knowledge graphs and large language models in neuro-symbolic systems. The E-KELL framework is a prominent example, where a knowledge graph provides structured context that guides the output of a language model [9]. This reduces hallucinations and improves explainability by grounding outputs in structured data. However, current implementations remain limited. Many rely on static data, do not incorporate real-time updates and lack explicit modeling of geospatial risk. Most importantly, they are not integrated into operational alerting infrastructures and therefore do not produce personalized, actionable alerts for end users.

Across all these areas, a consistent pattern emerges: strong individual components, but weak integration. Data integration, risk assessment and language generation are each well developed, yet they are rarely combined into a single system [34, 42]. Existing approaches fail to connect data, location and decision-making in a way that enables dynamic, context-aware and personalized alerts.

Research Gap. This highlights a clear research gap. There is currently no system that integrates knowledge graphs as a structured data backbone, risk assessment for dynamic spatial segmentation and large language models for personalized communication into a unified architecture. As a result, the core problems of current systems remain unsolved: alerts stay generic and imprecise, irrelevant notifications cause alert fatigue, and user trust declines. Addressing this gap requires an end-to-end system that connects these components in order to reduce irrelevant alerts, increase spatial precision, and deliver personalized, context-aware warnings that strengthen user trust.

3 Ambition Beyond the State of the Art

The analysis of the current state of the art reveals a fundamental limitation across existing approaches: while significant progress has been made in individual components such as data integration, event detection and natural language processing, these elements remain largely disconnected. Current emergency alert systems lack the ability to combine heterogeneous data sources, contextual understanding, and personalized decision-making into a unified and operational system [30, 34, 37].

This project advances the state of the art in three key dimensions: (1) integration of heterogeneous data into a unified semantic structure, (2) dynamic and user-specific risk evaluation, and (3) grounded generation of personalized emergency communication.




Fig. 1. Proposed system architecture integrating heterogeneous data sources into a knowledge graph, enabling downstream risk assessment and personalized alert generation.

Data Integration. Relevant data sources in this domain include meteorological APIs (the Austrian national weather and geophysics service, formerly ZAMG), IoT and sensor networks, satellite data, and geospatial infrastructure datasets such as OpenStreetMap [2, 8]. These are typically accessed via REST-based APIs or streaming pipelines (e.g. event-driven ingestion using frameworks such as Apache Kafka) [36]. Existing systems often rely on static datasets or loosely coupled pipelines, which limits interoperability and temporal consistency [34]. The proposed system introduces a unified ingestion layer that integrates real-time and static data into a coherent representation, enabling consistent downstream processing and improving situational awareness.

Knowledge Graph. The system leverages knowledge graph technology to represent entities such as hazards, locations, infrastructure, and populations, including their spatial and temporal relationships. The implementation follows RDF-based graph models combined with ontology-driven design, enabling semantic interoperability and structured querying via SPARQL [21, 26].

Alternative approaches such as relational databases or NoSQL systems provide scalability but are not well suited for modeling complex cross-domain relationships [42]. Property graph databases (e.g. Neo4j) offer flexibility but lack standardized semantics and interoperability across domains. RDF-based knowledge graphs therefore provide a suitable foundation for integrating heterogeneous disaster-related data into a consistent and interoperable structure [34, 38].

Risk Assessment. Risk modeling is based on established geospatial analysis techniques such as spatial overlays, proximity analysis, and threshold-based classification of hazard intensity [23, 33]. These methods can be implemented using GIS frameworks (e.g. QGIS, PostGIS) and rule-based decision logic [2, 18].

While machine learning models may improve predictive accuracy, they typically operate as black boxes: they cannot justify why a specific user was assigned a given risk level, which makes their outputs difficult to audit and to defend in safety-critical decisions [1]. The proposed approach therefore emphasizes rule-based and explainable risk computation, where every classification can be traced back to explicit criteria, for example a combination of hazard intensity thresholds, proximity, and infrastructure exposure. This transparency directly supports the human oversight and traceability required for high-risk AI applications. The key advancement lies in transforming static, region-based classifications into dynamic, fine-grained risk zones that are continuously updated with incoming hazard data and linked to individual user context, enabling more precise and accountable decision-making.

LLM-based Communication. Recent advances in natural language processing enable the use of large language models for context-aware communication. A suitable approach in this domain is retrieval-augmented generation (RAG), where structured information is retrieved from the knowledge graph and risk layer and incorporated into controlled prompts [19, 28].

Purely generative approaches without structured grounding are unsuitable due to hallucination risks and lack of traceability [5]. Instead, the system constrains outputs through structured inputs. Model choices remain flexible and may include locally hosted models (e.g. Mistral Large 3, Llama 4) or privacy-compliant cloud-based solutions, depending on deployment requirements [4, 27]. This ensures both data sovereignty and reliable communication in safety-critical environments.

Delivery. Alert dissemination builds on existing telecommunication infrastructures such as cell broadcast systems [13, 43]. These systems ensure scalability and reliability but typically operate on coarse regional targeting. Advances in mobile network technologies (4G/5G) enable more precise approximation of user location [11, 45], allowing alerts to be aligned with dynamic risk zones. The proposed system enhances these mechanisms without replacing them, maintaining robustness while improving spatial precision and relevance.

Overall Contribution. The key contribution lies in the integration of these technologies into a coherent architecture. While individual components such as semantic data models, geospatial risk analysis, and language generation are well established, they are rarely combined into a unified system [34, 42]. By linking structured data, spatial reasoning, and grounded communication, the proposed approach enables a transition from generic broadcast alerts to personalized, context-aware emergency communication. This directly addresses the core limitations identified in current emergency alert systems and provides a scalable and operational solution for real-world deployment. Ultimately, it aims to reduce irrelevant alerts and alert fatigue while strengthening user trust and responsiveness.

From Fragmented Research to an End-to-End System. The proposed architecture is designed explicitly around the limitations identified in the previous section. Existing knowledge graph approaches provide structured data integration, but they do not connect this structured knowledge to user-level risk and alert delivery. Geospatial risk modelling enables hazard-zone classification, but usually remains disconnected from personalized communication. LLM-based approaches support natural language generation, yet without structured grounding they are not sufficiently reliable in safety-critical settings. Even emerging neuro-symbolic approaches that combine knowledge graphs and language models still rely mostly on static data and do not integrate real-time geospatial risk evaluation into operational alerting infrastructures. AURORA addresses these gaps by connecting these previously isolated capabilities in a single end-to-end workflow, thereby creating the basis for the methodological design and evaluation described in the following section.

4 Methodology

This project follows a Design Science Research (DSR) approach, which is particularly suitable for developing and evaluating innovative IT artifacts [16, 44]. The objective is to design, implement and assess a functional system that addresses a clearly identified problem.

Design Science Research follows an iterative process consisting of the phases design, build, evaluate and refine [16]. These phases are systematically applied to ensure both scientific rigor and practical relevance.

In the design phase, system requirements are derived from the identified limitations of current emergency alert systems. Based on these requirements, the conceptual architecture (Figure 1) is developed, defining the interaction between data integration, knowledge representation, risk assessment and communication. Requirements are refined through stakeholder engagement with public authorities and telecommunication actors to ensure alignment with real-world emergency management needs [31].

The build phase implements the system across multiple work packages. A knowledge graph is developed to integrate heterogeneous data sources (WP2), followed by a risk assessment component that derives dynamic and user-specific risk levels (WP3). Based on this structured input, a grounded language generation component is developed to produce context-aware emergency alerts (WP4). All components are implemented as part of an integrated architecture rather than as isolated modules.

In methodological terms, the project does not only develop individual technical artifacts, but deliberately tests whether the integration of these artifacts resolves the shortcomings identified in the state of the art. More specifically, the methodology is designed to evaluate whether structured data integration improves situational consistency, whether explainable geospatial risk assessment improves the precision of affected-user identification, and whether grounded language generation improves the relevance and reliability of alerts compared to generic broadcast communication. This makes the evaluation logic directly traceable to the research gap and ensures that the artifact is assessed not only in technical isolation, but also in relation to the concrete weaknesses of existing approaches. This also reflects the iterative logic of Design Science Research: requirements derived from current limitations are translated into artifact design, implemented as an integrated prototype, and then evaluated against the practical deficiencies of existing emergency alert systems.

The evaluation phase assesses both individual components and the integrated system. The LLM module is evaluated with respect to correctness, clarity and reliability, while the overall system is tested using realistic disaster scenarios such as floods or storms (WP5). Performance is measured using criteria such as accuracy of risk classification, relevance of alerts and reduction of unnecessary notifications compared to baseline systems such as AT-Alert [31, 37].

The iterative nature of DSR is reflected in continuous refinement cycles, where evaluation results and stakeholder feedback are used to improve system components and their integration [16].

The work packages are directly aligned with the DSR process. WP2 corresponds to the design of the data foundation, WP3 and WP4 to the build phase, and WP5 to system evaluation, while WP1 supports all phases through coordination and dissemination. This ensures a structured and traceable development from conceptual design to a validated prototype.

5 Workplan

WP 1 Project Management & Dissemination	Duration M1–M18	Lead Ministry	Participating Organisations + PMs TU Wien (10 PM), Utimaco (4 PM), A1 (10 PM), Ministry (15 PM) Total: 39 PM	Goal Ensure effective coordination, stakeholder engagement, and dissemination to support timely and aligned project execution.	Deliverables D1.1 Project plan & requirements – M2 – A1 D1.2 Dissemination report – M12 – Ministry D1.3 Final project report – M18 – Ministry
	Tasks T1.1 Coordination & reporting (M1–M18) Plan and coordinate project activities across all partners, including scheduling meetings, tracking progress against milestones, and preparing formal reports (project plan, mid-term, final). Owner: A1		T1.2 Stakeholder engagement (M2–M17) Identify and engage relevant stakeholders (e.g. civil protection authorities, telecom actors) through workshops and structured interactions to gather requirements and validate results. Owner: Ministry		T1.3 Dissemination activities (M6–M18) Disseminate project outcomes through scientific publications, industry channels, and stakeholder-oriented communication formats. Owner: TU Wien
WP 2 Knowledge Graph & Data Foundation	Duration M1–M7	Lead TU Wien	Participating Organisations + PMs TU Wien (21 PM), Utimaco (8 PM), A1 (4 PM), Ministry (7 PM) Total: 40 PM	Goal Develop a structured knowledge graph integrating heterogeneous data sources to represent disaster-related context.	Deliverables D2.1 Ontology & data model – M4 – TU Wien D2.2 Knowledge graph prototype – M7 – TU Wien
	Tasks T2.1 Data source identification (M1–M2) Identify and analyse relevant data sources (e.g. weather data, geospatial data, infrastructure data) and define requirements for integration into the system. Owner: TU Wien		T2.2 Ontology design (M2–M4) Develop an ontology defining key entities (hazards, locations, infrastructure) and relationships, enabling structured representation and reasoning. Owner: TU Wien	T2.3 Knowledge graph implementation (M3–M6) Implement the knowledge graph using appropriate technologies (e.g. RDF/SPARQL), including data ingestion and storage mechanisms. Owner: TU Wien	T2.4 Interface design (M5–M7) Design and implement interfaces (APIs) to enable integration of the knowledge graph with downstream components (risk model and LLM module). Owner: Utimaco
WP 3 Risk Assessment & Geospatial Analysis	Duration M4–M10	Lead Utimaco	Participating Organisations + PMs TU Wien (18 PM), Utimaco (20 PM), A1 (10 PM), Ministry (7 PM) Total: 55 PM	Goal Design and implement a geospatial risk assessment model to identify and classify dynamic hazard zones.	Deliverables D3.1 Risk model specification – M7 – TU Wien D3.2 Risk assessment prototype – M10 – Utimaco
	Tasks T3.1 Geospatial modelling (M4–M6) Develop geospatial models to represent hazard propagation and spatial relationships between affected areas and infrastructure. Owner: TU Wien		T3.2 Risk logic development (M5–M7) Define and implement risk calculation logic, including thresholds and classification into risk levels (e.g. low, medium, high). Owner: TU Wien	T3.3 Real-time data integration (M7–M9) Integrate dynamic or simulated real-time data streams to enable continuous updating of risk assessments. Owner: Utimaco	T3.4 Risk zone generation (M9–M10) Generate and visualize dynamic risk zones based on model outputs for use in downstream alert generation. Owner: TU Wien
WP 4 Grounded LLM Communication	Duration M7–M12	Lead TU Wien	Participating Organisations + PMs TU Wien (24 PM), Utimaco (18 PM), A1 (5 PM), Ministry (8 PM) Total: 55 PM	Goal Develop a grounded LLM-based system for generating context-aware and reliable emergency alerts.	Deliverables D4.1 LLM communication module – M11 – TU Wien D4.2 Evaluation report – M12 – TU Wien
	Tasks T4.1 RAG architecture design (M7–M8) Design a grounded generation pipeline that combines structured inputs from the knowledge graph and risk model with LLM-based text generation. Owner: TU Wien		T4.2 Prompt engineering (M8–M10) Develop and refine prompt templates to ensure consistent, structured, and context-aware emergency alert generation. Owner: TU Wien	T4.3 Integration with KG and risk outputs (M10–M11) Integrate structured data inputs (knowledge graph + risk zones) into the LLM pipeline to constrain and guide generated outputs. Owner: Utimaco	T4.4 Evaluation of outputs (M11–M12) Evaluate generated alerts with respect to correctness, clarity, and hallucination reduction using defined evaluation criteria. Owner: TU Wien
WP 5 Integration, Validation & Evaluation	Duration M10–M18	Lead A1	Participating Organisations + PMs TU Wien (19 PM), Utimaco (14 PM), A1 (21 PM), Ministry (15 PM) Total: 69 PM	Goal Integrate system components and validate the prototype in realistic disaster scenarios.	Deliverables D5.1 Integrated prototype – M14 – A1 D5.2 Validation report – M16 – Ministry D5.3 Final evaluation report – M18 – TU Wien
	Tasks T5.1 System integration (M10–M13) Integrate all system components (knowledge graph, risk model, LLM module) into a unified prototype architecture. Owner: Utimaco		T5.2 Prototype deployment (M13–M14) Deploy the prototype in a controlled test environment enabling end-to-end functionality. Owner: Utimaco	T5.3 Scenario testing (M14–M16) Define and execute realistic disaster scenarios (e.g. flood, storm) to test system behaviour and performance. Owner: Ministry	T5.4 Evaluation against baseline systems (M16–M18) Evaluate system performance against reference compared to existing approaches (e.g. AT-Alert baseline). Owner: TU Wien

Table 1. Overview of the five work packages, including structure, tasks, responsibilities and deliverables across the project lifecycle.

The project is structured into five interdependent work packages (WPs) that follow the overall system architecture from data integration to validation. Each work package corresponds to a key system component and reflects the sequential development process defined in the methodology.




Table 2. Gantt Chart of the AURORA Project Work Packages

Table 1 provides a structured overview of all work packages, including their duration, responsibilities, core tasks, and expected deliverables. The visual structure mirrors the system architecture by organizing the development process from data integration (WP2), over risk assessment (WP3), to communication (WP4) and final system validation (WP5), while WP1 ensures coordination and stakeholder alignment throughout the project.

The work packages are tightly interconnected. WP2 establishes the knowledge graph as the data backbone of the system, which enables the geospatial risk modelling in WP3. These results are then translated into context-aware alerts in WP4 through the grounded LLM approach. WP5 integrates all components into a unified prototype and evaluates the system under realistic disaster scenarios, ensuring both technical validity and practical applicability. WP1 runs across the entire project, ensuring coordination, stakeholder involvement, and dissemination of results.

Together, the work plan reflects a structured yet iterative development process, where each work package builds upon the previous one while contributing to the overall system functionality.

The project is structured over the period 2027–2028 and follows a phased but overlapping implementation logic. As shown in the Gantt chart 2, project management and dissemination (WP1) run throughout the full duration of the project, ensuring continuous coordination, stakeholder involvement, and reporting.

The technical work progresses from foundational components toward integrated system validation. WP2 starts early in the project in order to establish the data and knowledge foundation on which the later technical components depend. Based on this groundwork, WP3 and WP4 are developed in parallel, allowing risk assessment and communication components to evolve iteratively alongside one another. In the later stages of the project, WP5 brings these components together in an integrated prototype and provides sufficient time for validation, scenario-based testing, and final evaluation.

Overall, the timeline reflects a realistic development strategy in which dependencies between work packages are respected, while overlap between phases enables iterative refinement and efficient use of project time.

6 Budget

Category	Cost (€)
Personnel	1,431,155
Travel	30,047
Equipment / Infrastructure	230,540
Other (dissemination, workshops)	57,505
Total (planned)	1,749,247
Contingency	221,531
Total budget	1,970,778

Table 3. Total project budget allocation.

The total project budget amounts to € 1,970,778, of which € 1,749,247 are allocated to planned project activities and € 221,531 are reserved as contingency. The overall budget is based on a bottom-up costing approach and remains below the maximum funding ceiling.

Cost Structure. Personnel costs amount to € 1,431,155 and are based on 258 person-months across the four partners. An average monthly personnel rate was calculated for each partner based on the roles they provide and then multiplied by the amount of person-months they contribute to the project. The FWF-standard was used for official project personnel rates. For other roles, Austrian salary levels were researched and increased by 20% to reflect employer-side personnel costs.

The travel costs of € 30,047 were calculated based on the planned project meeting and dissemination structure and include six plenary meetings, including meeting-related lunch and refreshments. The travel budget also includes participation in EU dissemination events as well as further travel for stakeholder engagement and validation activities. Equipment / infrastructure costs amount to € 230,540. These cover the project-specific technical resources required for implementation and validation, including compute resources for geospatial and LLM-based processing, hosting and staging environments, data storage and backup, geospatial services and APIs, and technical monitoring and test environments. Standard office equipment is not included.

Other direct costs amount to € 57,505 and include dissemination, outreach, publication, and communication activities. These cover the project website and visual identity, open-access publication fees, policy briefs and factsheets, dissemination materials, final stakeholder-oriented communication activities, and editing, translation, and printing support.

The contingency of € 221,531 is retained at project level as a reserve for uncertainties in implementation, integration, and validation. The contingency corresponds to approximately 12.7% of the planned project budget, which represents a moderate reserve level for an innovative multi-partner research and prototype project such as this one. Further details on the distribution of costs by partner and by work package are provided in the appendix.

Overall Allocation. The distribution of the planned budget across partners reflects both the allocated person-months and the specific technical and organisational roles within the project. Higher personnel shares are associated with

technically intensive components such as knowledge graph development, risk modelling, and system integration, while coordination, dissemination, and validation activities are reflected in the allocation of travel and additional costs. Similarly, the distribution of costs across work packages follows a mixed bottom-up logic. Personnel costs are directly derived from the underlying person-month allocations, while travel, infrastructure, and other direct costs are assigned based on the functional requirements of each work package. Technically intensive work packages therefore carry a higher share of infrastructure and implementation-related costs, whereas coordination and dissemination activities are reflected in work packages focused on project management and stakeholder engagement. Overall, the budget reflects a realistic and activity-based allocation of resources that is closely aligned with the technical architecture and work plan of the project.

7 Management Structure

The project follows a structured management approach that ensures alignment between the technical system architecture and the organizational execution across partners. The management structure is designed to support coordination across interdependent work packages while enabling efficient task-level execution.

General Assembly. The General Assembly consists of representatives from all four project partners and serves as the strategic decision-making body of the project. It convenes quarterly for regular progress reviews, additionally at each major project milestone for deliverable validation, and on an ad-hoc basis when urgent decisions are required. Meetings focus on reviewing cross-work-package progress, validating deliverables, and resolving dependencies between components. All decisions are documented centrally to ensure transparency and traceability throughout the project lifecycle.

Project Coordinator. The project is coordinated by A1, which provides an experienced project manager to serve as the central interface between all partners, the funding body, and external stakeholders. The coordinator is responsible for formal reporting and compliance, monitoring overall project progress against milestones and budget, and ensuring alignment across all work packages. Beyond administrative coordination, the coordinator acts as the first escalation point for inter-partner conflicts and manages external communication with civil protection authorities and other relevant stakeholders.

Technical coordination is led by TU Wien, which assumes responsibility for overseeing the scientific and implementation-related aspects of the project. Given its expertise in knowledge graphs, semantic technologies, and AI-based system design, TU Wien ensures consistency across the technical work packages and provides guidance on architectural decisions and system integration throughout the development process.

Work Package Leaders. Work package leaders meet weekly to report on the current status of their respective work packages, flag emerging risks or blockers, and coordinate dependencies with other work packages. These meetings ensure that intermediate outputs remain compatible across components and that technical interfaces are aligned before handover between work packages. Each work package leader is accountable to the General Assembly and the coordinator for the timely delivery of their work package's deliverables.

Task-Level Execution. Task leaders are responsible for the day-to-day operational execution within their assigned tasks. They coordinate small interdisciplinary teams, distribute work, track progress, and escalate issues to the work package leader when necessary. A key responsibility of task leaders is translating high-level conceptual designs into implementable components and ensuring consistency across system layers.

Collaboration Tools and Workflow. The project relies on integrated digital collaboration workflows. GitHub is used for version control, code management, and documentation of technical decisions. Confluence serves as the central

platform for documentation and knowledge sharing across partners, and additionally provides built-in Kanban boards for task tracking and sprint planning at the task level. This combination allows for both synchronous coordination and asynchronous collaboration across partners.

External Advisor. An external advisor provides targeted feedback at key project milestones, covering expertise in civil protection, telecommunications, and AI ethics to ensure the project remains aligned with operational, technical, and societal expectations. Given that AURORA employs AI in a safety-critical context, ethical guidance is a particularly important aspect of this advisory role. , a leading expert in AI ethics, has been engaged to provide guidance on the ethical design and implementation of the AI components. She will be consulted selectively at relevant project stages rather than on a regular basis, contributing targeted input where ethical considerations are most critical, particularly in the design of the risk classification logic and the grounded LLM communication layer.

8 Procedures

Risk Management The project follows a structured risk management approach based on PRINCE2 principles, adapted to the scale and context of a multi-partner research and prototype project. Risks are derived from partner expertise and discussed collectively at the project outset, with the coordinator and work package leaders jointly deciding on probability, impact, and appropriate response strategies. All identified risks are categorised into five domains: Technical, Operational, Integration, Stakeholder, and Project Management and documented in a central risk register maintained by the coordinator (Table 4).

Each risk is assessed in terms of probability and impact, resulting in an overall severity rating (High, Medium, or Low), complemented by a proximity assessment indicating when the risk is most likely to materialise. Based on this assessment, a response strategy and a designated Risk Owner are assigned, ensuring clear accountability for mitigation and monitoring. Beyond threats, the project also monitors emerging opportunities identified during execution, which are assessed using the same process and integrated into the system where relevant.

Risk reviews are conducted on a weekly basis within work package leader meetings, allowing emerging risks to be identified and addressed early. Should a risk be classified as High, an immediate coordination loop is triggered between the relevant work package leaders and the project coordinator, with escalation to the General Assembly if no resolution is reached at that level. The risk register is reviewed comprehensively on a monthly basis, ensuring that risk management remains a continuous and integrated part of project governance throughout the full project lifecycle.

Innovation Management. Innovation is managed through an iterative development approach aligned with DesignScience Research [16], emphasizing continuous refinement of system components based on evaluation results and stakeholder feedback. Technical development is organized using GitHub, ensuring transparent version control, reproducibility, and collaborative development across all partners.

Intellectual property is managed through a Consortium Agreement based on the DESCA (Development of a Simplified Consortium Agreement) framework [14], which serves as the contractual foundation governing ownership, access rights, and exploitation of results among all four project partners.

The project follows a selective open-source strategy, distinguishing between components that can be made publicly available and those that must remain restricted due to their integration into critical infrastructure. The system architecture and overall layer setup, the ontology and data model developed in WP2, the risk assessment module, and the pipeline from risk assessment output to grounded LLM communication will be released as open source. Components directly tied to operational telecommunication infrastructure will not be publicly released.

ID	Risk Author	Date Registered	Risk Category	Risk Description	Probability x Impact	Proximity	Response Category	Status	Risk Owner
1	TU Wien	06/2026	Technical	Heterogeneous data sources cannot be integrated consistently into the knowledge graph.	High	Early project	Reduce	Active	TU Wien
2	Utimaco	05/2026	Technical	Relevant data is delayed or incomplete, reducing the quality of risk assessment.	High	Mid project	Reduce	Active	Utimaco
3	TU Wien	05/2026	Operational / Technical	Risk classification may assign users to incorrect risk levels.	High	Mid project	Reduce	Active	TU Wien / Utimaco
4	TU Wien	05/2026	AI / Communication	The LLM may generate unclear, misleading, or hallucinated emergency alerts.	High	Mid to late project	Reduce	Active	TU Wien
5	Utimaco	06/2026	Integration	Single components may be difficult to integrate into one prototype.	Medium	Late project	Reduce	Active	Utimaco / AI
6	Ministry	03/2026	Project Management	Delays in one work package may affect dependent work packages.	Medium	Through-out	Reduce	Active	Ministry
7	Ministry	03/2026	Stakeholder	Stakeholder requirements may be unclear or change during the project.	Medium	Through-out	Reduce	Active	Ministry
8	AI	05/2026	Validation	Scenario testing may not sufficiently reflect realistic emergency situations.	Medium	Late project	Reduce	Active	Ministry / AI

Risk / Impact
High Risk > 7.500 €
Medium > 1.000 €
Low Risk < 1.000 €

Table 4. Risk register

Project results will be disseminated through multiple channels. Findings will be submitted to the ISCRAM conference (Information Systems for Crisis Response and Management), the leading international venue for research at the intersection of information systems and crisis management. A technical report documenting the system architecture and key design decisions will be published as a preprint. Additionally, the project will leverage the Ministry’s existing relationships with public authorities and industry stakeholders to present results at relevant national and European events.

Data Management Data management is guided by the FAIR principles and is closely aligned with the knowledge graph-based system architecture of the project.

The project integrates data from multiple heterogeneous sources. Meteorological and environmental data is sourced from GeoSphere Austria (formerly ZAMG), accessed via their public Data Hub API. Geospatial infrastructure data is obtained from OpenStreetMap under the Open Database License (ODbL). IoT sensor and network data is provided by A1 from infrastructure in proximity to mobile network towers. Environmental monitoring data is obtained in collaboration with the Austrian Federal Ministry of Climate Action and Environment. All data sources are used in compliance with their respective licenses.

These heterogeneous inputs are integrated into a unified semantic representation within the knowledge graph, ensuring interoperability across all system components. Personal data is not directly stored; location-related information is processed in aggregated or approximate form to ensure privacy and regulatory compliance in accordance with GDPR. With respect to the FAIR principles, the project applies them as follows. Project outputs including the ontology, knowledge graph, and source code are published via Zenodo, ensuring persistent identification through DOIs and long-term findability beyond the project lifetime. Access to outputs follows a differentiated strategy: components released as open source are freely accessible, while data directly tied to A1 telecommunication infrastructure remains restricted due to operational and security constraints. Interoperability is ensured through the use of established Semantic Web standards, in particular RDF and SPARQL, which enable integration with external systems and datasets. To support reusability, all open-source components are released under explicit licenses – Apache 2.0 for software and CC BY 4.0

for the ontology and data model accompanied by technical documentation sufficient for independent reuse and further development.

Legal and Ethical Considerations. The project addresses legal and ethical aspects as an integral part of system design, ensuring compliance with relevant European and Austrian legislation throughout development and deployment. Regarding data protection, the system is developed in compliance with the General Data Protection Regulation (GDPR, EU 2016/679). Location-related information is processed in aggregated form only and no personal data is directly stored. As AURORA employs AI in a safety-critical context, the project is aligned with the EU Artificial Intelligence Act (EU AI Act, Regulation EU 2024/1689). Emergency alert systems involving automated risk classification may be considered high-risk AI applications under this regulation. The project therefore emphasizes explainability, human oversight, and traceability as core design requirements. Alert dissemination via cell broadcast is developed in alignment with the European Electronic Communications Code (EECC, Directive EU 2018/1972), and the system design is compatible with the EU Civil Protection Mechanism (Decision No 1313/2013/EU). Ethical considerations include the risk of incorrect alerts, biases in risk classification, and impacts on public trust. These are mitigated through rule-based decision models, grounded language generation, and human oversight for critical scenarios.

9 Project Partner Organisations

The project combines academic, industrial, and public-sector expertise, ensuring that the technical system design is grounded in both scientific rigor and real-world applicability.

TU Wien is a leading technical university with strong expertise in knowledge graphs, artificial intelligence, and data-driven systems. Its research focus on semantic technologies and system design provides the foundation for structuring and integrating heterogeneous data, as well as for developing and evaluating advanced system architectures. Beyond its scientific contribution, TU Wien assumes technical coordination of the project, ensuring consistency across all technical work packages and guiding key architectural decisions throughout development.

Utimaco brings extensive experience in security-critical systems and large-scale infrastructure solutions. Beyond its expertise in system integration, backend architectures, and secure data handling, Utimaco has direct experience in the development of emergency alert and disaster warning systems in Austria, making it a uniquely qualified partner for the implementation of robust and reliable alerting components in high-stakes environments.

A1 contributes deep expertise in telecommunications infrastructure and large-scale communication systems. Its experience with mobile networks and broadcast technologies provides essential insights into real-world deployment constraints and enables the alignment of system outputs with existing communication infrastructures. A1 is an essential partner for the project, as emergency alerts are disseminated directly through the A1 mobile network via cell broadcast technology, making their infrastructure the backbone of the final delivery layer. Beyond the technical layer, A1 also leads project coordination and brings significant experience in managing large-scale multi-partner projects, ensuring structured and efficient execution across all work packages.

The Austrian Ministry of Internal Affairs contributes domain expertise in civil protection, emergency management, and regulatory frameworks. Its role ensures that the system design remains aligned with operational requirements, public warning systems, and stakeholder expectations within real-world crisis scenarios. The Ministry is an essential partner as the operational end user (Bedarfsträger) and the authority responsible for civil protection and public warning in Austria, providing the regulatory alignment, operational validation, and formal endorsement required for real-world deployment. Furthermore, the development and deployment of emergency alert systems in Austria requires regulatory

alignment with civil protection authorities and access and compliance that would not be achievable without the Ministry’s involvement and formal endorsement.

10 Expected Outcomes and Impact

Technical Outcomes. The project will deliver a validated, end-to-end prototype of a context-aware emergency alert system integrating an RDF-based knowledge graph, explainable geospatial risk assessment, and grounded language generation. It is demonstrated in at least three realistic scenarios (flood, severe thunderstorm, and wildfire) and validated against a reconstructed AT-Alert baseline, meeting the project objectives: at least 95% correct answers on 20 competency questions, risk-zone updates within two minutes, and improved precision and recall of affected-user identification with fewer irrelevant alerts than the broadcast baseline. The system architecture, the emergency-risk ontology and data model, the risk-assessment module, and the risk-to-LLM pipeline are released as open artifacts on Zenodo (Apache 2.0 and CC BY 4.0), with the ontology aligned to the Common Alerting Protocol (CAP) for interoperability with existing warning infrastructures.

Academic Impact. The project targets at least three peer-reviewed publications: on the emergency knowledge graph and ontology (targeting ESWC or ISWC), on the explainable geospatial risk model evaluated against expert-annotated ground truth (targeting the International Journal of Disaster Risk Reduction), and on the grounded LLM layer and its factual grounding against the knowledge graph (targeting an NLP and crisis-informatics venue). A consolidated system paper will be submitted to the ISCRAM conference, complemented by a technical report published as a preprint. Together with the openly released artifacts, these results form the basis for a Horizon Europe follow-up within Cluster 3 (Disaster-Resilient Society, topics DRS-01 and DRS-02) and a national scale-up within the FFG KIRAS security-research programme.

Industrial Impact. For Utimaco, which already develops emergency alert systems in Austria, the validated prototype forms the basis for a future commercial product in safety-critical alerting. For A1, the project demonstrates risk-aware, real-time alerting over its existing cell-broadcast infrastructure, deepens its expertise in IoT-based data integration, and showcases its infrastructure capabilities for stakeholder communication.

Societal Impact. AURORA directly addresses over-alerting and alert fatigue in current Austrian warning systems. By replacing identical broadcasts with risk-based, context-aware alerts, it is expected to measurably reduce irrelevant notifications and strengthen user trust, assessed through stakeholder workshops and user-acceptance studies. This is achieved without compromising data protection, as no personal data is stored, location is processed in aggregated form, and the system complies with the GDPR and the high-risk requirements of the EU AI Act through explainability and human oversight. For the Ministry and civil protection authorities, the findings are translated into operational recommendations via a policy brief, strengthening national preparedness.

Sustainability and Outlook. Beyond the funded period, each result has a clear owner: Utimaco advances the prototype toward a commercial alerting product, A1 carries the delivery layer into network operations, and the Ministry anchors the system in civil-protection practice, while the open-sourced ontology and pipeline remain available to the research community via Zenodo. This aligns AURORA with the EU Preparedness Union Strategy and the Union Civil Protection Mechanism as a transferable building block for the next generation of European public warning. Ultimately, AURORA turns generic broadcast alerts into precise, explainable, and trusted communication that reaches the right people at the right moment.

References

- [1] Ahmed Saeed Ali Rashed Aladawi and Ahmad Nur Aizat Ahmad. 2023. A study of factors influencing the adoption of artificial intelligence in crisis management. *International Journal of Sustainable Construction Engineering and Technology* 14, 5 (2023), 416–425.
- [2] Raffaele Albano and Aurelia Sole. 2018. Geospatial methods and tools for natural risk management and communications. 470 pages.
- [3] Amin Anjomshoaa, Hannah Schuster, Johannes Wachs, and Axel Polleres. 2023. Towards Crisis Response and Intervention Using Knowledge Graphs - CRISP Case Study. In *Advanced Information Systems Engineering Workshops* (1 ed.) (*Lecture Notes in Business Information Processing*), Marcela Ruiz and Pnina Soffer (Eds.). Springer, Germany, 67–73. doi:10.1007/978-3-031-34985-0_7
- [4] Manuela Benary, Xing David Wang, Max Schmidt, Dominik Soll, Georg Hilfenhaus, Mani Nassir, Christian Sigler, Maren Knödler, Ulrich Keller, Dieter Beule, et al. 2023. Leveraging large language models for decision support in personalized oncology. *JAMA Network Open* 6, 11 (2023), e2343689.
- [5] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*. 610–623.
- [6] Grégoire Burel and Harith Alani. 2018. Crisis event extraction service (crees)-automatic detection and classification of crisis-related content on social media. (2018).
- [7] Jared J Cash. 2009. Alert fatigue. *American Journal of Health-System Pharmacy* 66, 23 (2009), 2098–2101.
- [8] Petros Chavula, Fredrick Kayusi, Gilbert Lungu, and Agnes Uwimbabazi. 2025. The current landscape of early warning systems and traditional approaches to disaster detection. *LatIA* 3 (2025), 77.
- [9] Minze Chen, Zhenxiang Tao, Weitong Tang, Tingxin Qin, Rui Yang, and Chunli Zhu. 2024. Enhancing emergency decision-making with knowledge graphs and large language models. *International Journal of Disaster Risk Reduction* 113 (2024), 104804.
- [10] Marco Esteban Silva Chuquillanqui, Alejandro Jose Llanos Garcia, Ronald Paucar Curasma, and Daniel Diaz Ataucuri. 2015. Study of emergency warning broadcasting systems. In *2015 IEEE International Symposium on Broadband Multimedia Systems and Broadcasting*. IEEE, 1–6.
- [11] Andrea Conti, Flavio Morselli, Zhenyu Liu, Stefania Bartoletti, Santiago Mazuelas, William C Lindsey, and Moe Z Win. 2021. Location awareness in beyond 5G networks. *IEEE Communications Magazine* 59, 11 (2021), 22–27.
- [12] Julie Bu Daher, Tom Huygue, Nathalie Hernandez, and Patricia Stolf. 2022. Inferring New Information from a Knowledge Graph in Crisis Management: A Case Study.. In *KDIR*. 43–54.
- [13] Pankaj Kumar Dalela, Ankita Saldhi, Pramod Bhawe, and Vipin Tyagi. 2020. Common alerting protocol compliant emergency warning and alert system for legacy broadcasting networks. In *2020 IEEE Conference on Cognitive and Computational Aspects of Situation Management (CogSIMA)*. IEEE, 1–5.
- [14] DESCA Core Group. 2024. DESCA HE 2.0 – Model Consortium Agreement for Horizon Europe. <https://www.desca-agreement.eu>. Coordinated by Fraunhofer and the Helmholtz Association.
- [15] Giorgia Di Capua and Stefan Rahmstorf. 2023. Extreme weather in a changing climate. *Environmental Research Letters* 18, 10 (2023), 102001.
- [16] Aline Dresch, Daniel Pacheco Lacerda, and José Antônio Valle Antunes Jr. 2014. Design science research. In *Design science research: A method for science and technology advancement*. Springer, 67–102.
- [17] Iryna Dymun, Oleh Basytiuk, and Andriy Dumyn. 2024. Graph-Based Approaches for Multimodal Medical Data Processing.. In *IDDm*. 337–348.
- [18] Valentina Ferretti and Gilberto Montibeller. 2019. An integrated framework for environmental multi-impact spatial risk analysis. *Risk analysis* 39, 1 (2019), 257–273.
- [19] Premanand Ghadekar, Shreyash Tekade, Dhawal Sakharwade, Sayee Zanzane, Ankur Tripathi, and Shivam Tiwadi. 2025. Real-time crisis response optimization with Retrieval Augmented Generation (RAG). In *Intelligent Computing and Communication Techniques*. CRC Press, 86–91.
- [20] Ala Harika, Gunapriya Balan, H Pal Thethi, Ajay Rana, K Varada Rajkumar, and Mustafa Abdulhussein Al-Allak. 2024. Harnessing the power of artificial intelligence for disaster response and crisis management. In *2024 International Conference on Communication, Computer Sciences and Engineering (IC3SE)*. IEEE, 1237–1243.
- [21] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard De Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, et al. 2021. Knowledge graphs. *ACM Computing Surveys (Csur)* 54, 4 (2021), 1–37.
- [22] Muhammad Imran, Carlos Castillo, Fernando Diaz, and Sarah Vieweg. 2015. Processing social media messages in mass emergency: A survey. *ACM computing surveys (CSUR)* 47, 4 (2015), 1–38.
- [23] Ratiranjana Jena, Biswajeet Pradhan, Ghassan Beydoun, Abdullah Al-Amri, and Hizir Sofyan. 2020. Seismic hazard and risk assessment: a review of state-of-the-art traditional and GIS models. *Arabian Journal of Geosciences* 13, 2 (2020), 50.
- [24] Shaheen Khatoun, Majed A Alshamari, Amna Asif, Md Maruf Hasan, Sherif Abdou, Khaled Mostafa Elsayed, and Mohsen Rashwan. 2021. Development of social media analytics system for emergency event detection and crisis management. *Computers, Materials, & Continua* 68, 3 (2021), 3079.
- [25] Grace Kim, Andrea Martel, David Eisenman, Michael Prellip, Armen Arevian, Kerri L Johnson, and Deborah Glik. 2019. Wireless Emergency Alert messages: Influences on protective action behaviour. *Journal of Contingencies and Crisis Management* 27, 4 (2019), 374–386.
- [26] Ora Lassila, James Hendler, and Tim Berners-Lee. 2001. The semantic web. *Scientific american* 284, 5 (2001), 34–43.
- [27] Joel Lawson, Viktor Krasnoshchok, Jian Sun, Danielle J Williams, Marco Campion, and Roberto Giacobazzi. 2026. Improving Crisis Communication and Public Policy Decisions Through Real-Time LLM Data Summarization. (2026).

- [28] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems* 33 (2020), 9459–9474.
- [29] Sara K McBride, Danielle F Sumy, Andrea L Llenos, Grace A Parker, Jeffrey McGuire, Jessie K Saunders, Men-Andrin Meier, Pascal Schuback, Douglas Given, and Robert de Groot. 2023. Latency and geofence testing of wireless emergency alerts intended for the ShakeAlert® earthquake early warning system for the West Coast of the United States of America. *Safety science* 157 (2023), 105898.
- [30] Dennis S Mileti and John H Sorensen. 1990. *Communication of emergency public warnings: A social science perspective and state-of-the-art assessment*. Technical Report. Oak Ridge National Lab., TN (USA).
- [31] National Academies of Sciences, Medicine, Division on Engineering, Physical Sciences, Computer Science, Telecommunications Board, Committee on the Future of Emergency Alert, Warning Systems, and Research Directions. 2018. Emergency alert and warning systems: Current knowledge and future research directions. (2018).
- [32] Haneesha Perusomula, Vamsi Marriwada, Sai Krishna Vallepu, Karthikeya Sesham, Roshan Chowdary Gudapati, and Jyothiraditya Garikipati. 2023. Geo-Fencing and Overspeed Alert SMS System. In *2023 Second International Conference on Electronics and Renewable Systems (ICEARS)*. IEEE, 584–589.
- [33] Biswajeet Pradhan, Pravat Kumar Shit, Gouri Sankar Bhunia, Partha Pratim Adhikary, and Hamid Reza Pourghasemi. 2022. *Spatial Modelling of Flood Risk and Flood Hazards*. Springer.
- [34] Hemant Purohit, Rajaraman Kanagasabai, and Nikhil Deshpande. 2019. Towards next generation knowledge graphs for disaster management. In *2019 IEEE 13th international conference on semantic computing (ICSC)*. IEEE, 474–477.
- [35] Jacipt Alexander Ramón-Valencia, Jordi Rafael Palacios-González, German Ricardo Santos-Granados, and Jarol Derley Ramón-Valencia. 2019. Early warning system on extreme weather events for disaster risk reduction. *Revista Facultad de Ingeniería Universidad de Antioquia* 92 (2019), 80–87.
- [36] Theofanis P Raptis and Andrea Passarella. 2023. A survey on networked data streaming with apache kafka. *IEEE access* 11 (2023), 85333–85350.
- [37] Abdul-Akeem Sadiq, Ratna B Dougherty, Jenna Tyler, and Rebecca Entress. 2023. Public alert and warning system literature review in the USA: Identifying research gaps and lessons for practice. *Natural hazards* 117, 2 (2023), 1711–1744.
- [38] Hannah Schuster, Amin Anjomshoaa, and Axel Polleres. 2026. Building Incident-Centric Knowledge Graphs: Enhancing Crisis Management with Data-Driven Insights. In *Dynamics of Disasters: Hybrid Threats*. Springer, 101–111.
- [39] Jeannette Sutton and Michele M Wood. 2025. Opting Out: Over-Alerting and Warning Fatigue in the Era of Wireless Emergency Alerts. *Journal of Contingencies and Crisis Management* 33, 3 (2025), e70076.
- [40] Alexandros Vassiliades, Grigorios Stathopoulos-Kampilis, Gerasimos Antzoulatos, Spyridon Symeonidis, Sotiris Diplaris, Stefanos Vrochidis, Nick Bassiliades, and Ioannis Kompatsiaris. 2024. XR4DRAMA a knowledge-based system for disaster management and media planning. *The Knowledge Engineering Review* 39 (2024), e1.
- [41] Sarah Vieweg, Amanda L Hughes, Kate Starbird, and Leysia Palen. 2010. Microblogging during two natural hazards events: what twitter may contribute to situational awareness. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1079–1088.
- [42] Xinzhi Wang, Mengyue Li, Weiwang Chen, Yige Yao, Zhennan Li, Yi Liu, and Hui Zhang. 2024. Review and Application of Knowledge Graph in Crisis Management. *International Journal of Software Engineering and Knowledge Engineering* 34, 03 (2024), 393–425.
- [43] Yanfeng Wang, Dazhi He, Lianghui Ding, Wenjun Zhang, Wei Li, Yiyang Wu, Ning Liu, and Yao Wang. 2017. Media transmission by cooperation of cellular network and broadcasting network. *IEEE Transactions on Broadcasting* 63, 3 (2017), 571–576.
- [44] Ningshuang Zeng, Luxuan Han, Yan Liu, Jingfeng Yuan, and Qiming Li. 2025. Design science research (DSR) in construction: Theoretical conceptualization of practice and practical realization of theory. *Automation in Construction* 176 (2025), 106298.
- [45] Ping Zhang, Jian Lu, Yan Wang, and Qiao Wang. 2017. Cooperative localization in 5G networks: A survey. *Ict Express* 3, 1 (2017), 27–32.

11 Appendix

Partner	PM	Personnel (€)	Additional (€)	Total (€)
TU Wien	92	471,690	44,220	515,910
Utimaco	64	401,434	65,427	466,861
A1	50	309,475	72,940	382,415
Ministry	52	248,556	135,505	384,061
Total	258	1,431,155	318,092	1,749,247

Table 5. Budget distribution across project partners.

The distribution of the planned budget across partners reflects both the allocated person-months and the specific technical and organisational roles of each partner in the project. TU Wien has the largest personnel budget due to its strong research contribution across the technical work packages and its role in scientific coordination. Utimaco carries a high budget share because of its central role in technical development, implementation, and risk-related system components. A1 contributes strongly to integration and deployment-related activities, while also supporting overall management by providing a dedicated project manager. The Ministry contributes to project coordination, stakeholder engagement, regulatory and legal aspects, operational validation, and dissemination-related activities. The additional cost column includes each partner’s share of travel, infrastructure, and other direct project costs. These costs were allocated according to the expected functional contribution of each partner rather than by simple proportional distribution.

Work Package	PM	Cost (€)
WP1	39	251,114
WP2	40	287,916
WP3	55	380,576
WP4	55	368,300
WP5	69	461,341
Total	258	1,749,247

Table 6. Cost distribution across work packages.

The budget per work package was calculated using a mixed bottom-up approach. Personnel costs were derived directly from the revised person-month allocation by multiplying the PM contribution of each partner in each work package by the corresponding partner-specific average monthly rate. Travel, infrastructure, and other direct costs were then allocated according to the actual purpose of each work package rather than through a purely proportional split.

WP1 includes project management, stakeholder engagement, dissemination planning, and coordination activities across the full project duration. It therefore carries a comparatively high share of travel and outreach-related costs, but only a limited share of technical infrastructure costs.

WP2 focuses on the structured data foundation of the project, including ontology design, knowledge graph implementation, and interface design. It therefore carries a significant share of technical infrastructure and system-related development costs, while dissemination and travel costs remain comparatively limited.

WP3 covers the development of the geospatial risk assessment and dynamic hazard modelling components. As a technically intensive work package, it includes a high personnel share and a substantial share of infrastructure costs related to processing, modelling, and technical implementation.

WP4 focuses on grounded LLM-based communication, including prompt design, integration of structured inputs, and output evaluation. This justifies a high personnel and infrastructure share, particularly for technically demanding development and testing activities.

WP5 includes system integration, deployment in a controlled prototype environment, scenario-based validation, and final evaluation. It therefore represents the highest-cost work package, combining strong personnel effort with integration resources, validation-related travel, and dissemination-related project outputs.