AI-based concepts for Crisis Propagation Forecasting and Early Warning in Urban Areas

Andrea Tundis
Institute for the Protection of Terrestrial Infrastructures
German Aerospace Center (DLR)
Rathausallee 12, 53757, Sankt Augustin, Germany
andrea.tundis@dlr.de

Maximilian Hummel
Institute for the Protection of Terrestrial Infrastructures
German Aerospace Center (DLR)
Rathausallee 12, 53757, Sankt Augustin, Germany
maximilian.hummel@dlr.de

Jonas Gunkel
Institute for the Protection of Terrestrial Infrastructures
German Aerospace Center (DLR)
Rathausallee 12, 53757, Sankt Augustin, Germany
jonas.gunkel@dlr.de

Claudio Savaglio
Department of Informatics, Modeling, Electronics and Systems
University of Calabria (UNICAL)
Rende, Italy
csavaglio@dimes.unical.it

***Abstract*—By 2050, over 68% of the global population is expected to live in cities, increasing demands on energy grids, transportation networks, and public services. This urban growth will intensify the complexity of infrastructure systems and heighten vulnerability to natural disasters, cyber-physical failures, and crises. In urban emergency management, unexpected events frequently occur that traditional radar and weather systems fail to detect. Additionally, existing warning mechanisms are often disconnected from predictive models, causing delays in issuing context-aware alerts. Current systems also struggle to deliver tailored warnings that account for the dynamic circumstances of individual citizens. In this large panorama, AI-powered technologies, particularly those using machine learning and natural language processing, offer promising solutions to overcome these limitations and improve urban crisis management. Based on that, this paper proposes an integrated crisis warning system to address these challenges. It presents an AI-centered pipeline that connects an imputation model, a forecasting model, and a crisis advisor, creating a robust and adaptive warning system. The proposal is evaluated through a scenario-based case study conducted in Darmstadt (Germany). Storm scenarios are used to test its forecasting and alert-generation capabilities, highlighting its potential to enhance urban resilience and citizen safety.**

***Index Terms*—Crisis Forecasting, Early Warning, Safety, Urban Critical Infrastructures, Machine Learning, Artificial Intelligence, Large Language Model.**

I. Introduction

Modern urban environments face challenges due to rapid population growth, more frequent natural disasters [1], and the increasing complexity of infrastructure systems. Traditional crisis management frameworks, reliant on manual decision-making, and delayed information, are increasingly inadequate in this dynamic scenario. For instance, during the 2021 floods in Germany, delayed warnings resulted in catastrophic human and economic losses [2]. Such and other incidents underscore the need for crisis warning systems capable of preemptive forecasting, and context-aware communication.

Indeed, in managing emergencies in cities, unexpected events can happen that are not detected by traditional equipment, on the other hand they can still be detected by other sources such as local sensors and even through reports from people. However, such sources do not cover all areas of a city, by creating gaps and uncertainty. Since crises can escalate across both space and time, this incomplete information makes it difficult to generate accurate short-term predictions and timely warnings for affected citizens. In fact, warning mechanisms are often separated from predictive models, leading to delays in generating context-aware alerts due to dependency on external forecasting services. They also do not provide personalized information that reflect the changing situations of individual citizens. Artifical Intelligece (AI)-driven technologies can help bridge existing gaps by enabling continuous urban monitoring, spatio-temporal crisis modeling, and automated behavioral guidance [3]. However, current implementations often favor isolated technical solutions over integrated approaches, missing opportunities for comprehensive crisis management. For instance, machine learning (ML) models can leverage diverse data sources—from sensor readings to real-time citizen reports—to predict crisis trajectories with spatio-temporal precision. Combined with natural language processing systems, these predictions can be transformed into personalized warnings, actionable recommendations, and interactive support tailored to individual needs. This integrated approach aligns with directives, like those from the United Nations Office for Disaster Risk Reduction [4].

In this context, the paper deals with the delivering of short-term, timely crisis warnings by integrating live sensor

The research activities related to this work are being conducted in the context of ”urbanModel” and ”MoDa (Models and Data for Future Mobility Supporting Services)” projects that are funded by the German Aerospace Center (DLR) e.V. This activity has been carried out in cooperation with the LOEWE Zentrum emergenCITY.

2325-2944/25/$31.00 ©2025 IEEE
DOI 10.1109/DCOSS-IoT65416.2025.00122

data and citizen reports into a ML-based crisis warning system. In particular, the paper presents a set of AI-based concepts as a comprehensive pipeline that links an imputation model, a forecasting model, and a crisis advisor to form a robust crisis warning system, thus exemplifing distributed collective intelligence. The proposal is evaluated through a scenario-based case study in the urban area of Darmstadt, (Germany), with storm scenarios serving both for forecasting and warning generation in compliance to the German guidelines.

The rest of the paper is organized as follows: Section II provides an overview of the related work; whereas the proposed concepts and related architectural model are elaborated in Section III. Implementation details and the obtained results are discussed in Section IV. Section V concludes the paper and highlights the further works.

II. Background on Warning Message Systems

This section surveys various approaches for stakeholder warning and communication mechanisms, starting with an examination of traditional warning systems, followed by infrastructure-based solutions, emerging social media and AI-driven approaches. For each approach a discussion of the current limitations is contextually provided.

a) Traditional Warning Systems: Traditional warning systems rely on push notifications, Short Message Service (SMS), or app-based alerts to inform the public about potential crises. These systems aggregate data from authoritative sources such as weather services, disaster control agencies, and municipal authorities to disseminate relevant warnings. Typically, they allow users to customize alerts based on location or crisis type, making them adaptable to individual needs. The goal is to provide real-time, actionable information to enhance public safety during emergencies. Applications like NINA [5], KATWARN [6], and BIWAPP [7] exemplify national and regional warning systems that integrate data from local and federal authorities. These systems distribute alerts via multiple channels, such as smartphone apps, SMS, and emails. Similarly, apps like "Mein DRK" and "Sicher Reisen" focus on specific contexts, such as first aid and international travel safety [8]. On a global scale, platforms like Disaster Alert monitor multiple hazards and provide location-specific warnings. Traditional warning systems provide alerts and basic instructions to large populations, leveraging established communication channels. *Limitations* — They are scalable and accessible, requiring minimal user interaction or technical expertise. However, these systems exhibit weaknesses of limited personalization, dependence on timely authoritative sources, and lack of open source software.

b) Infrastructure-Based Systems: Infrastructure-based systems rely on existing telecommunication networks or specialized organizational infrastructures to deliver crisis alerts on a broad scale. They typically operate independently of user-installed apps, focusing instead on sending messages directly through channels like mobile network cells or organizational IT platforms. This enables rapid, large-scale dissemination of warnings—even when conventional communication lines are overloaded—by leveraging built-in functionalities of mobile networks or enterprise systems. *Limitations* — Infrastructure-based systems can rapidly reach wide audiences, but the approach has some limitations. Cell Broadcast excels at delivering large-scale alerts instantly yet cannot tailor messages to individual users, since all compatible devices in a targeted zone receive the same notification [9]. Meanwhile, sistems like safeREACH focuses on internal organizational contexts, which limits its applicability beyond those structured environments, reducing reach and flexibility for the public [10].

c) Social Media- and AI-Driven Approaches: Social media-driven approaches leverage user-generated data and real-time communication on social platforms to detect, classify, and respond to emerging threats. Such systems draw on the large user base and instant feedback loop offered by social media, enabling rapid identification of crisis hotspots and direct interaction with stakeholders. Social media-based warning systems benefit from the broad coverage and scale of their user communities, which can span entire countries or even transcend national borders. Additionally, these platforms enable two-way communication, allowing both crisis response entities and the public to share critical information in real time. This approach often involves proprietary or closed-source implementations (e.g., Facebook Safety Check [11]) but can also include open-source AI frameworks (e.g., LLAMA2 [12]) that allow for ongoing community-driven improvements. *Limitations* — Although social media-driven methods exploit massive user networks and continuous content streams for rapid crisis identification, they often depend on proprietary platforms that require active user accounts, and the reliability of crowdsourced reports can suffer when misinformation circulates. Additionally, LLM-based solutions face the "hallucination" problem, where the model may generate inaccurate or misleading data. Such false information is a serious vulnerability in crisis management.

A. BBK Recommendations for Warning Messages

As discussed above, it emerges that warning messages play a central role in crisis management, and their effectiveness is essential to enabling timely and provide appropriate responses from affected populations. The clarity, structure, and credibility of these messages significantly influence public perception and compliance. Miscommunication or ambiguity in warnings can lead to confusion or delays in response. To ensure consistency and efficacy, the German *Federal Office of Civil Protection and Disaster Assistence* (Bundesamt für Bevölkerungsschutz und Katastrophenhilfe–BBK) has provided structured guidelines for formulating warnings. These guidelines aim to minimize misinterpretation, and align warning dissemination with established crisis communication

TABLE I
DESCRIPTION OF THE BBK WARNING MESSAGE RECOMMENDATIONS

Recommendation
Issuers/Senders: Clearly identify the issuer to enhance credibility
Target Groups: Tailor warnings to specific target audiences
Crisis Identification: Clearly and comprehensibly describe the crisis
Action Recommendations: Provide clear instructions tailored to the crisis
Consistency: Formulate warnings consistently and without contradictions
References to Additional Information: Provide links or references to additional resources
Updates: Regularly update warnings to maintain relevance
All-Clear Messages: Clearly indicate when a crisis has subsided
Text Structure: Structure warning texts logically to improve readability
Readability: Ensure the text is understandable for a broad audience
Relevance of the Warning: Highlight the severity and urgency of the crisis
Text Length: Keep the message short and concise without omitting essential details
Precision: Formulate warnings precisely and unambiguously
Objectivity: Write neutrally without exaggeration
Sentence Construction: Use simple and active sentences
Word Choice: Use commonly understood terms and avoid or explain technical jargon
Spelling and Grammar: Avoid errors to ensure credibility and comprehension
Transparency: Provide transparent information and clearly state any uncertainties
Highlighting: Emphasize the most important information
Warning Words: Use specific warning words within the text

principles. The German BBK guidelines [13], which have been shortly reported in Table I, outline key considerations for constructing effective warning messages, and will be taken into account within our proposed solution.

III. Architectural Model and AI-based Concepts

This section introduces the proposed architectural model, that is depicted in Fig. 1. It represents the fundamentals of the envisioned AI-based system, that has been structured into the four main building blocks discussed hereinafter: *(a)* Input source *(b)* Data Preparation Layer, *(c)* Crisis Forecasting Layer and *(d)* Stakeholder Alert and Communication Layer.

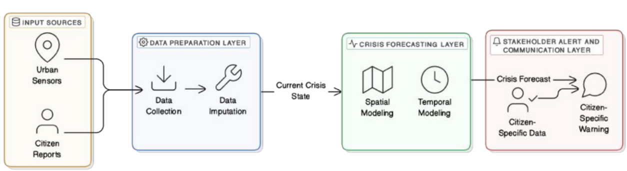

Fig. 1. Warning pipeline and building blocks of the proposed architecture

a) Input Sources: The system under consideration is designed for urban areas, which are geographically defined using boundaries specified by latitude and longitude coordinates. This urban area serves as the spatial domain within which all crisis modeling and management activities are performed. The urban area is discretized into a grid of $n \times n$ cells, where each cell corresponds to a unique spatial region (see Fig. 2).

The choice of n allows for flexible spatial resolutions. This discretization ensures a clear spatial assignment, as each spatial event or entity can be mapped to a specific cell. Each grid cell g_{ij} is defined as a polygonal region

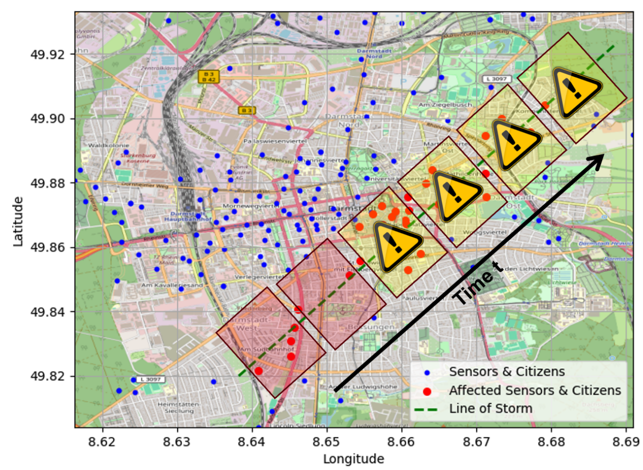

Fig. 2. Grid-based approach: affected entities along the linear disaster path

with direct adjacency relationships to its neighboring cells. This relationship allows the modeling of spatial propagation dynamics between adjacent cells. A grid cell at a specific time may contain various sensors and citizens, thus providing two different input sources: *(i) Urban Sensors* from Critical Infrastructures (CIs): each grid cell may host one or more sensors, monitoring infrastructure elements such as water pumps or electrical transformers. Please note that different sensor types can have different attributes, e.g., a water pump sensor has two attributes "water level" and "flow rate", defining its state; *(ii) Citizens Reports*: for each citizen contained in a certain time in a given cell, it is possible to define its multi-attribute status in analogy with the urban sensors. For example, consider a citizen with two attributes indicating whether they are outside and whether their home location is affected by a crisis.

Additionally, a citizen's status may include an attribute for reporting crises in real time. This attribute is called ”Crisis Report”, and allows citizens to transmit information about the crisis. This report is modeled as an additional attribute specifying ”Crisis Type” (e.g. Flood, Fire, Earthquake, Storm, Other), ”Location”, ”TimeOfIncident”, optional ”Description”, ”Severity” into a 1-5 range, optional ”ImmediateNeeds” specifications (e.g. Evacuation, Medical Help, Shelter, Other).

b) Data Preparation Layer: It is the first operative stage of the crisis warning system. Its primary role is to collect and to aggregate input source data about *Urban Sensors* and *Citizen Reports* at each discrete time t to form the source grid. For each cell, a preliminary binary state is determined using a predefined function based on available sensor and citizen reports. Uncertainty arises for cells where no sensor or citizen data is available. To handle this uncertainty, an *imputation model* is applied that leverages spatial relationships within the gride cells to compute a probability indicating the likelihood that cell is affected. These probabilities are then converted into binary states by applying a certain threshold. Importantly, if a cell was already affected in the original grid, this status is retained in the imputed grid.

Hence, the resulting imputed grid comprises all cells that were already affected in the original grid as well as those cells indicated by the imputation model. This process ensures that the imputed grid reflects the crisis state seamlessly, by incorporating both observed data and spatially inferred effects.

c) Crisis Forecasting Layer: It follows the Data Preparation Layer and represents the second stage of the warning pipeline. Its input is the imputed grid for each istant, which represents the crisis propagation up to that point. Over time, a sequence of imputed grids is accumulated, forming a historical context. This sequence encapsulates the temporal evolution of the crisis and serves as input to the spatiotemporal forecasting model, which then predicts future crisis propagation. Therefore, the model forecasts grids indicating, for each grid cell, whether the cell will be affected in the next istant. These probability values are subsequently thresholded to obtain binary outputs, thereby mimicking the forecasted crisis propagation.

The model is applied in an autoregressive manner to predict multiple future timesteps. After producing a forecast, the newly predicted grid is appended to the context sequence. Therefore the context is expanded with forecasted data which is then used as input for predicting the subsequent. This iterative approach is repeated until the desired number of future timesteps has been generated (see Fig. 3).

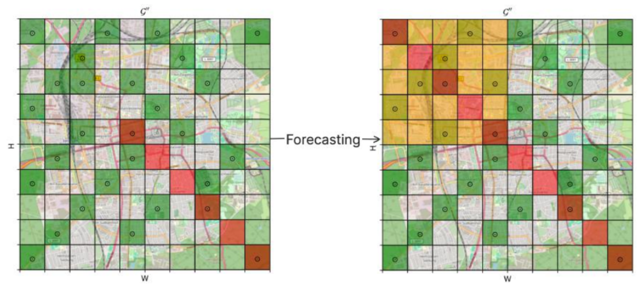

Fig. 3. Forecasted Crisis - Yellow color highlights urban areas which could be affected and impacted from the crisis propagation. Red color highlights urban areas which could be affected and strongly impacted from the crisis propagation.

d) Stakeholder Alert and Communication Layer: it is the final stage of the proposed pipeline. It receives the predicted grids for each future timestep, generated by the previous layer. In addition to these forecasted grids, this layer integrates the current ”CitizenStatus”. Using this contextual information, the system formulates an input for a *crisis advisor* that can assist citizens by providing situation-specific guidance based on: *(1)* the most recent citizen status; *(2)* the forecasted crisis propagation, i.e., which cells are likely to be affected soon; and *(3)* the verified government knowledge sources, specifically government documents.

Relying on official government documents ensure that all guidance adhere to established protocols. The crisis advisor acts as an interactive assistant for citizens with two tasks:

1) *Citizen-Specific Warnings*: generate situation-aware warnings in natural language. These warnings include clear instructions on how to respond to the unfolding crisis, given the citizen's current ”CitizenStatus” and the predicted crisis progression.
2) *Interactive Chat Interface*: provide a conversational platform for citizens seeking additional information, clarifications, or assistance during a crisis. This chat interface may be accessed after receiving a warning for follow-up questions or proactively at any stage of the crisis, for example during the crisis preparation phase.

The process of creating a citizen-specific warning begins by setting up a dedicated directory that stores official documents from the BBK, which serve as the foundational knowledge base for the crisis advisor. Each document in this directory is assigned metadata. This metadata includes information about the crisis stage addressed by the document (e.g., preparation or response phases) and a list of keywords summarizing its content. For instance, documents might be tagged with keywords such as “flood”, “evacuation”, or “shelter” to group relevant materials.

When a citizen's status is received along with a forecast of the crisis, additional details are added to describe their situation in relation to the predicted event. For instance, if the system expects the citizen to be in an area that will be affected, their status is updated to indicate that they are at risk. Additionally, information about the direction of the crisis in relation to their location is also recorded.

The final step is the generation of warnings based on three key inputs: the ”extended citizen status”, the ”predicted crisis grid”, and the ”verified knowledge base” derived from government documents. Using this data, a language model (LM) should produce specific warnings in natural language. These warnings include actionable instructions customized to the situation, such as evacuation orders, recommended shelter locations, or other crisis-specific guidance. The generated warnings should adhere to the official government requirements for generating textual warnings of Table I: therefore, the crisis advisor's output should follow the format below to meet the guidelines for warning creation.

As depicted in Fig. 4, this warning design meets the “Crisis Identification” guideline (Guideline 3) by explicitly describing the threat. This warning aligns with official requirements, which recommend a concise description of the crisis to help recipients comprehend the situation quickly. The text also fulfills the “Action Recommendations” guideline (Guideline 4) by offering instructions on seeking shelter and adopting protective measures. Moreover, the structured format uses distinct sections for Notification Type, Date, Time, Warning, and Behavioral Advice, meeting the “Text Structure” guideline (Guideline 9) through clear text organization. In summary, this layer completes the end-to-end pipeline by providing crisis information to citizens, ensuring they can take appropriate actions before a crisis fully impacts their location. Moreover, this layer offers a textual crisis advisor, which is available before, during, and after a crisis to meet citizens' specific needs and situational context.

IV. Implementation and Qualitative Results

The final layer in the presented warning pipeline orchestrates the transformation of forecasted crisis grids into targeted citizen alerts that integrate crisis forecasts, citizen statuses, and verified government information. Additionally, this layer provides a chat interface for citizens to ask questions proactively, for example, during the crisis preparation phase. A RAG approach is implemented to achieve this functionality, as it bridges the gap between static knowledge bases and dynamic generation by retrieving and grounding outputs in verified and contextually appropriate information, reducing the risk of hallucination in LLMs [14]. It constructs a knowledge base from official government documents stored locally, retrieves context-relevant passages using a vector index, and then fuses these references with citizen-specific details to generate crisis notifications or guidance. In the context of crisis communication, where the dissemination of misinformation can have severe consequences, RAG proves particularly advantageous by incorporating official documents and directives into its retrieval process. This process ensures that all generated alerts and advisories are rooted in authoritative sources, aligning with the need for adherence to official guidelines.

a) Document/Input Preprocessing RAG: Before using the RAG approach, documents are selected to establish a knowledge base with crisis relevant documents. The BBK provides freely accessible information on how to behave before, during, and after a crisis. This information is used and collectively saved as documents within a directory, enabling the index creation for the RAG. In this paper, 36 crisis documents were manually selected containing general crisis advisories, emergency tips and crisis-specific information, e.g., on how to behave within a storm. All the documents are official government information from the BBK [13]. A vector-based indexing technique is utilized, transforming documents into embeddings. As the documents can become large, they are chunked in smaller pieces of a maximum of 1024 tokens. This token size results from the usage of the default tokenizer ”cl100k” from tiktoken [15] and the default chunk size within ”llama index” [16]. This chunk size and tokenizer is used as a starting point and can be adapted through the llama index library if it is required. Too large chunks result in too much information in the embedding, while too small chunks loose the surrounding context. This vector-based index captures semantic relationships between documents, enabling retrieval. Upon initialization, there is a check for a previously persisted index of government documents. If none is found, a new index is created by reading all documents within the specified directory, assigning metadata such as crisis stage (preparedness, response, or recovery) to each document. This metadata is useful for filtering content to only the most important segments. The metadata is assigned to each document through an automated classification process. For this purpose, a LLM gets document titles and the content as input, to categorize each document in the corresponding crisis stage (preparedness, response, recovery). This automatic metadata assignment using LLMs was implemented to ensure seamless integration of new documents in the system, which normally would require manual metadata assignment for each document. Moreover, the metadata includes keywords of the document which are extracted using keyword extractor in llama index. Once the index is loaded, the RAG approach can be used to retrieve relevant and official crisis information using a LLM.

In Summary, this RAG approach combines forecasted grids, citizen statuses, and official government directives into natural-language warnings. Moreover, citizens can establish a chat, to inform themselves about crisis relevant information.

b) Evaluation of the RAG-Based Approach: The evaluation of the RAG-based approach is conducted based

⚠ STORM WARNING

Type of notification: Storm Warning **Date/Time:** January 25, 2025, 11:48 PM

Directly Affected: Yes

Warning: A severe storm is approaching your area. Expect strong winds and heavy rain.

Behavioral advice:

- Stay indoors and keep away from windows.
- Avoid rooms that could be damaged by falling trees.
- Turn on the radio or television to receive updates.
- In an emergency, call the fire department at 112.

Note: This automated notification was generated using LLMs. Verify critical information through official channels.

Fig. 4. Design and Storm Warning generated by the RAG-based approach

on three distinct scenarios: A citywide blackout, a storm in the city, and a citizen preparing for a potential upcoming crisis. Within each of these scenarios, three different citizen situations are defined, representing different contexts of the citizen. For each citizen situation in every scenario, the RAG-based approach is queried five times to generate relevant information relating to the citizen's current situation. This evaluation approach ensures diversity across scenarios and user situations while accounting for the inherent non-deterministic nature of LLMs.

For the implementation, the GPT-4o model was accessed via the OpenAI API [17]. The retrieval step was executed with a ”temperature” setting of 0 to ensure deterministic behavior in crisis document retrieval. In contrast, the grounding step utilized a ”temperature” of 0.2 to enable a more flexible yet strictly compliant formulation of warnings while adhering to the constraints specified in the system prompt. Additionally, the embeddings required for the RAG process were generated using OpenAI embedding models.

A qualitative evaluation of the warnings generated for citizens was conducted. Since the RAG output undergoes a grounding step to enforce compliance with the required format for warnings, individual warnings are qualitatively assessed against the requirements outlined by the BBK of Table I. This step ensures that the generated warnings meet official standards and are suitable for practical application. More specifically, three warning messages are qualitatively evaluated based on official guidelines for warning message formulation provided by the BBK [13]. Table II presents an overview of the guideline criteria met or unmet for Warnings 1–3. Then, each warning is analyzed individually to highlight specific strengths and areas for improvement. The storm warning generated by the RAG system aligns with multiple criteria outlined in the BBK guidelines for effective crisis communication (see Fig. 4). The warning explicitly identifies the type of crisis (i.e. ”Storm Warning”) and specifies the date/time (i.e. January 25, 2025, 11:48 PM), fulfilling requirements under Crisis Identification (Guideline 3) and Precision (Guideline 13).

By stating that the user is directly affected, the message adheres to Target Groups (Guideline 2), as it personalizes the warning for citizens in the storm's path. The behavioral advice is a key strength, addressing Action Recommendations (Guideline 4) with specific, actionable steps. These instructions are tailored to mitigate risks from flying debris, structural damage, and glass shattering, reflecting Relevance of the Warning (Guideline 11). The inclusion of an emergency contact (”call 112”) and advice to monitor media further supports Consistency (Guideline 5). The Text Structure (Guideline 9) is logical, progressing from threat description to protective measures. Readability (Guideline 10) is ensured through simple, direct language (e.g. ”Expect strong winds”), avoiding technical jargon (Guideline 16). The warning's short length satisfies Text Length (Guideline 12), while its neutral tone aligns with Objectivity (Guideline 14). Moreover, references to Additional Information (Guideline 6) are also provided with the suggestion to turn on the radio or television.

Two limitations are notable. Firstly, the warning does not explicitly identify the issuer (Guideline 1), reducing transparency. While the note mentions automated generation via LLMs, it omits reference to an authoritative entity (e.g., meteorological service), which could affect credibility. Moreover,

TABLE II
EVALUATION OF THE THREE WARNINGS AGAINST 20 GUIDELINES

#	Guideline	At-Home Danger	Unprotected Outside	Driving Into Storm
1	Issuers/Senders	✗	✗	✗
2	Target Groups	✓	✓	✓
3	Crisis Identification	✓	✓	✓
4	Action Recommendations	✓	✓	✓
5	Consistency	✓	✓	✓
6	Additional Information	✓	✗	✗
7	Updates	✓	✓	✓
8	All-Clear Messages	✗	(✓)	(✓)
9	Text Structure	✓	✓	✓
10	Readability	✓	✓	✓
11	Relevance of the Warning	✓	✓	✓
12	Text Length	✓	✓	✓
13	Precision	✓	✓	✓
14	Objectivity	✓	✓	✓
15	Sentence Construction	✓	✓	✓
16	Word Choice	✓	✓	✓
17	Spelling and Grammar	✓	✓	✓
18	Transparency	✓	✓	✓
19	Highlighting	✓	✓	✓
20	Warning Words	✓	✓	✓

the warning is no All-Clear message as the citizen is directly affected within the current situation.

V. CONCLUSIONS AND FUTURE WORKS

The paper has dealt with early warning of crisis events in urban areas. In particular, a distributed collective intelligence system integrating heterogeneous data sources within an AI-based architectural model has been proposed, so as to enable a more effective and dynamic communication with the users, centered on they specific needs and provided information. A grid approach was employed to model the environment as well as evaluate the evolution of the reference crisis event, while a LLM was trained, using government documents and official guidelines, to support the natural language processing tasks and thus automate and dynamize the communication process with users. The model has been tested by simulating a crisis situation in the urban area of the city of Darmstadt. The results of the qualitative assessment have shown its capability to generate tailored information to help the citizens during the emergency based on their current state, while also respecting the guidelines and remaining compliant with government directives.

Future work includes *(i)* refining the concepts through formal notation, *(ii)* deepening the experimentation and address privacy of citizen-generated reports, as well as *(iii)* providing a quantitative evaluation of the system, also assessing the viability of lightweight AI alternatives (e.g., distilled LLMs).

REFERENCES

[1] Wetter extreme in Deutschland. 2023. Online available: 2025-03-24 - https://de.statista.com/statistik/studie/id/162237/dokument/wetterextreme-in-deutschland/

[2] A. Fekete and S.Sandholz. “Here Comes the Flood, but Not Failure? Lessons to Learn after the Heavy Rain and Pluvial Floods in Germany 2021”. In: Water 13.21 (2021), p. 3016. doi: 10.3390/w13213016.

[3] R. Casadei, G. Fortino, D. Pianini, W. Russo, C. Savaglio, and M. Viroli, ”A development approach for collective opportunistic edge-of-things services,” *Information Sciences*, vol. 498, pp. 154–169, 2019.

[4] United Nations Office for Disaster Risk Reduction. Global Assessment Report on Disaster Risk Reduction 2022: Our World at Risk: Transforming Governance for a Resilient Future. Geneva: United Nations Office for Disaster Risk Reduction, 2022

[5] Bundesamt für Bevölkerungsschutz und Katastrophenhilfe (BBK). Notfall-Informations- und NachrichtenApp (NINA). https://www.bbk.bund.de/DE/Warnung-Vorsorge/Warn-App-NINA/warnapp-nina_node.html Accessed: 2025-02-11

[6] Lazka - Python API Client for KATWARN. https://github.com/lazka/katwarn-api Accessed: 2025-02-11

[7] BIWAPP – Bürger Info & Warn App. https://www.biwapp.de/ Accessed: 2025-02-11

[8] Deutsches Rotes Kreuz. Mein DRK. https://mobil.drk-intern.de Accessed: 2025-02-11

[9] GSMA. Mobile Network Public Warning Systems and the Rise of Cell-Broadcast. https://www.gsma.com/mobilefordevelopment/wp-content/uploads/2013/01/Mobile-NetworkPublic-Warning-Systems-and-the-Rise-of-Cell-Broadcast.pdf Accessed: 2025-03-10

[10] safeREACH. – Alarmierungssoftware & Krisenmanagement für Unternehmen. https://safereach.com/de/ Accessed: 2025-03-10

[11] Facebook - Inc. Facebook Safety Check. https://www.facebook.com/about/safetycheck Accessed: 2024-12-21

[12] Hakan T. Otal and M. Abdullah Canbaz. “LLM-Assisted Crisis Management: Building Advanced LLM Platforms for Effective Emergency Response and Public Collaboration”, 2024, In: arXiv preprint arXiv:2402.10908

[13] Bundesamt für Bevölkerungsschutz und Katastrophenhilfe. Warnsysteme. Accessed: December 2024. - https://www.bmi.bund.de/DE/themen/themen-node.html

[14] O. Ayala and P. Bechard. “Reducing hallucination in structured outputs via RetrievalAugmented Generation”. In: Proc. of the 2024 Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Vol 6: Industry Track). Association for Computational Linguistics, 2024, pp. 228–238.

[15] OpenAI. tiktoken. https://github.com/openai/tiktoken. 2023. Accessed: 2025-02-02.

[16] Jerry Liu. LlamaIndex. Nov. 2022. doi: 10.5281/zenodo.1234. url: https://github.com/ jerryjliu/llama_index.

[17] OpenAI. OpenAI Platform Documentation. 2025. - https://platform.openai.com/docs