

ClimaEmpact: Domain-Aligned Small Language Models and Datasets for Extreme Weather Analytics

Deeksha Varshney^a, Keane Ong^a, Rui Mao^b,
Erik Cambria^b, Gianmarco Mengaldo^{*a}

^a*National University of Singapore, Singapore, Singapore*

^b*Nanyang Technological University, Singapore, Singapore*

Abstract

Accurate assessments of extreme weather events are vital for research and policy, yet localized and granular data remain scarce in many parts of the world. This data gap limits our ability to analyze potential outcomes and implications of extreme weather events, hindering effective decision-making. Large Language Models (LLMs) can process vast amounts of unstructured text data, extract meaningful insights, and generate detailed assessments by synthesizing information from multiple sources. Furthermore, LLMs can seamlessly transfer their general language understanding to smaller models, enabling these models to retain key knowledge while being fine-tuned for specific tasks. In this paper, we propose *Extreme Weather Reasoning-Aware Alignment (EWRA)*, a method that enhances small language models (SLMs) by incorporating structured reasoning paths derived from LLMs, and ExtremeWeatherNews, a large dataset of extreme weather event-related news articles. EWRA and ExtremeWeatherNews together form the overall framework, ClimaEmpact, that focuses on addressing three critical extreme-weather tasks: categorization of tangible vulnerabilities/impacts, topic labeling, and emotion analysis. By aligning SLMs with advanced reasoning strategies on ExtremeWeatherNews (and its derived dataset ExtremeAlign used specifically for SLM alignment), EWRA improves the SLMs' ability to

Email addresses: deeksha2@nus.edu.sg (Deeksha Varshney),
keane.ongweiyang@u.nus.edu (Keane Ong), rui.mao@ntu.edu.sg (Rui Mao),
cambria@ntu.edu.sg (Erik Cambria), mpegim@nus.edu.sg (Gianmarco Mengaldo*)

generate well-grounded and domain-specific responses for extreme weather analytics. Our results show that the approach proposed guides SLMs to output domain-aligned responses, surpassing the performance of task-specific models and offering enhanced real-world applicability for extreme weather analytics.

Keywords: Extreme Weather, Large Language Models, Small Language Models, Alignment, Reasoning

1. Introduction

The increasing occurrence and intensity of extreme weather events present significant dangers that affect both individuals and communities [1]. These events, including heatwaves, prolonged dry conditions, extreme precipitation, and tropical cyclones, have significant effects on human livelihoods and the natural environment, often leading to lasting and sometimes irreversible consequences. However, our limited understanding of their societal impacts hinders the development of effective disaster response and communication strategies [2, 3]. Enhancing access to accurate and comprehensive information is crucial to addressing this challenge. Such data can serve as a foundation for extreme weather analysis tasks that can provide crucial *extreme weather analytics* such as vulnerability and impact assessment [4, 5, 6, 7, 8, 9].

Large Language Models (LLMs) present a promising tool for providing comprehensive and accurate information on extreme weather events by efficiently extracting and synthesizing large climate-related data from various textual sources [10, 11, 12, 13]. Yet, there are several challenges on the pathway to achieve this goal, including outdated data that does not consider the latest events [14], and scarcity of expert-annotated data, which is critical for training and evaluation. These challenges can affect the accuracy of extreme weather analysis tasks by hindering the model’s ability to account for the latest extreme weather events, shifts in vulnerability, and evolving impact assessments. To mitigate these issues, some solutions have been proposed in the literature. For instance, ClimateGPT [15] implements various strategies, including a hierarchical retrieval augmentation strategy, drawing upon external knowledge sources like Wikipedia and IPCC AR6 reports. While integrating scientifically validated sources enhances the accuracy of LLM responses, it does not entirely eliminate the potential for generating inaccurate or misleading information [16].

Indeed, the critical gap posed by the lack of structured and up-to-date knowledge in the domain of extreme weather event analysis still persists, and requires urgent attention. Addressing this gap is crucial to comprehensively understand the evolving vulnerabilities and impacts associated with extreme weather events, encompassing both the tangible damages and the emotional consequences experienced by affected communities.

To fill this gap, we leverage the advanced reasoning capabilities of LLMs on news articles associated with extreme weather events. In particular, we prompt LLMs to generate detailed, step-by-step reasoning paths on three extreme-weather analysis tasks: (i) identification of vulnerability, impact, and emergency response of the region affected by the extreme event, (ii) topic/subtopic labeling and keyword extraction, and (iii) emotion analysis. These reasoning paths act as alignment data, providing clear, detailed explanations that illustrate how to approach and solve complex extreme weather analysis problems. For instance, when tasked with categorizing a sentence that describes a weather impact, the LLM does not simply provide the category but also explains the logical steps it followed, highlighting relevant keywords and contextual clues.

Building upon these alignment examples, we propose a novel domain-specific reasoning approach aimed at enhancing the analytical capabilities of language models for extreme weather analysis tasks. We introduce Extreme Weather Reasoning-Aware Alignment (EWRA), a fine-tuning method that transfers the advanced reasoning skills of LLMs to small language models (SLMs). EWRA trains SLMs using LLM-generated reasoning paths, enabling these models to decompose complex problems into manageable steps and provide clear, logical explanations, making them domain-specialized for providing extreme weather analytics. Our approach effectively applies a two-stage fine-tuning strategy: first, SLMs are trained with implicit reasoning to internalize reasoning logic instead of depending only on prompt patterns; second, the models are provided with detailed definitions of task categories within the prompts for explicit fine-tuning. This capability is particularly critical for retrieving useful extreme weather analytics, where robust representations capable of accurately identifying and characterizing extreme weather features and understanding the underlying rationale for predictions are as important as the predictions themselves [17, 18].

As part of this work, we also release a new dataset, ExtremeWeatherNews, which comprises news articles collected for 60 distinct extreme weather events. EWRA and ExtremeWeatherNews constitute the two elements of our Cli-

maEmpact framework (`climaempact.ai`) for near real-time extreme weather analytics. To evaluate the effectiveness of EWRA, we perform a series of experiments comparing its performance against baseline models on a specialized dataset derived from ExtremeWeatherNews, namely the ExtremeAlign dataset, prepared using LLM-generated reasoning paths. We evaluate multiple fine-tuning strategies, including standard and reasoning-based approaches, and show that EWRA delivers competitive performance compared to other fine-tuning methods across tasks, particularly on the Qwen2.5-3B-Instruct model. On the Vulnerability/Impact/Emergency assessment task, EWRA achieves 5.2% improvements in Spearman Rank Correlation over Reason Explicit-SFT, demonstrating stronger alignment with human-annotated reasoning patterns, while also showing comparable or superior performance on topic/subtopic labeling tasks and emotion analysis.

In summary, the key contributions of our work are as follows.

1. We introduce Extreme Weather Reasoning-Aware Alignment (EWRA), a novel approach for enhancing SLMs by transferring reasoning capabilities from LLMs through fine-tuning using synthetic alignment data for extreme weather event analysis.
2. We release ExtremeWeatherNews, a comprehensive collection of news articles related to 60 distinct extreme weather events, curated to support the study of various extreme weather related tasks. In order to implement EWRA, we also introduce a new alignment dataset called the ExtremeAlign, comprising rationales across three tasks.
3. We evaluate the performance of EWRA through a series of experiments using the ExtremeAlign dataset, demonstrating its ability to significantly improve domain-specific reasoning accuracy, as compared to task-specific baseline models.

The ClimaEmpact framework is deployed online for fast and near real-time analyses of extreme weather events.

2. Related Work

We focus on three key areas of research that are related to our, namely: standard NLP methods for extreme weather analysis, LLMs for extreme weather analysis, and SLMs alignment. Each of these areas and the related work is reported in the following.

Traditional NLP methods for extreme weather. Textual data, such as news reports, social media posts, and official bulletins, complement satellite-based datasets by providing context-rich and timely information on extreme weather impacts. Multiple natural language processing (NLP) approaches have utilized textual data to analyze and classify extreme weather impacts [19, 20, 21]. Social media data have also been exploited; for example, drought impacts in California have been categorized using BERT-based models applied to Twitter posts [22]. Similarly, automated text processing has mapped trajectories, impacts, and the aftermath of the 2021 European floods, providing broader categorization and insights into weather extremes [23].

Transformer-based models, such as ClimateBERT, have enhanced various NLP applications related to climate, including text classification, sentiment analysis, and fact-checking [24]. Moreover, affective computing and text mining techniques have proven valuable for capturing emotional nuances within textual datasets [25, 26]. An example includes analyzing public opinions on wildfires through a neurosymbolic approach employing BERT-based models [27]. Additionally, the ClimaText dataset has demonstrated the limitations of keyword-based approaches in climate discourse, emphasizing the advantages of context-aware methods like BERT for effective topic labeling [28, 29]. Recent studies have explored metaphorical concept mappings to analyze public perceptions of weather disasters, leveraging tools such as MetaPro to elucidate distinct public perspectives across various disaster types [2, 30].

LLMs for extreme weather. Current methods typically rely on task-specific models, limiting adaptability and efficiency. LLMs offer an integrated framework, leveraging unsupervised pretraining, instruction tuning, and reinforcement learning to enhance contextual understanding from extensive textual datasets [31]. These capabilities are increasingly utilized for disaster response tasks. Examples include FloodBrain, which automates flood impact reporting to expedite response times [32], and Llama2-based models designed for emergency classification and issuing public instructions during emergencies when 911 systems become overwhelmed [33]. Further, fine-tuned LLMs have improved disaster-related tweet classification for event identification and aid distribution [34], while other efforts have optimized LLM performance through prompt engineering techniques integrating textual and satellite data [35]. DisasterResponseGPT generates personalized disaster preparedness plans, enhancing response efficiency [36].

Several climate-specific LLMs have also emerged. ClimateGPT-2 fine-tunes GPT-2 with climate-specific data to support claim generation and fact-checking tasks [37]. The Arabic Mini-ClimateGPT focuses specifically on conversational Arabic climate instruction tuning using the Clima500-Instruct dataset [38], while ChatClimate [16] grounds climate-change-related responses in IPCC AR6 reports, enhancing accuracy and timeliness in climate conversations. Similarly, ClimateGPT synthesizes interdisciplinary climate knowledge through domain-specific instruction tuning and extensive climate-related resource training [15]. Despite their advantages, retraining and fine-tuning these LLMs are computationally intensive, leading to significant environmental impacts due to high carbon emissions.

Small Language Models (SLMs) alignment. Recent studies have focused on reasoning-based alignment methods to enhance small models using larger counterparts’ reasoning abilities. Reasoning alignment can significantly improve the performance of smaller models by transferring advanced reasoning from larger models [39]. Chain-of-Thought reasoning techniques have demonstrated substantial improvements in generalization capabilities [40], while other research has successfully transferred intermediate reasoning steps to enhance decision-making processes [41]. Knowledge Distillation (KD), a popular technique, reduces model complexity and latency while retaining accuracy [42, 43]. The Generalized Knowledge Distillation (GKD) framework extends traditional KD by incorporating structured knowledge to improve training outcomes [44].

Unlike conventional KD, our approach explicitly transfers intermediate reasoning structures rather than solely matching model outputs. Although SLM alignment has been explored extensively within general NLP contexts, its application specifically to extreme weather analysis remains understudied. To address this gap, we propose Extreme Weather Reasoning-Aware Alignment (EWRA), a method that aligns smaller models with larger LLMs by training them on reasoning paths generated by the latter. EWRA enables SLMs to decompose complex problems into structured reasoning steps and produce logical explanations, making them highly specialized for extreme weather analysis tasks. This approach offers computational efficiency advantages over traditional methods, substantially reducing the environmental footprint associated with large-scale model training. Additionally, we deliberately refrain from employing retrieval-augmented generation (RAG)-based systems [45, 16], as fine-tuning inherently incorporates real-time, domain-

specific information essential for handling nuanced queries related to extreme weather events thus limiting hallucination in response more effectively than retrieval-based approaches.

3. Methodology

In this section, we present the ClimaEmpact framework, composed of the ExtremeWeatherNews dataset (section 3.1), forming the foundation of our analysis; ExtremeAlign (section 3.2), a task-specific alignment dataset enabling SLMs to navigate complex reasoning tasks; and EWRA (section 3.3), a reasoning-aware fine-tuning method transferring reasoning capabilities from LLMs to smaller ones. EWRA improves SLMs’ ability to handle ambiguities and complexities inherent in extreme weather information.

3.1. *ExtremeWeatherNews (EWN) Dataset*

The ExtremeWeatherNews dataset comprises news articles collected for 60 distinct extreme weather events, and chosen based on ClimaMeter [46]¹. The extreme weather events include heatwaves, cold spells, extreme wind, and extreme precipitation; the name of each event considered is reported in Table A.1.

To build ExtremeWeatherNews, we employed a web scraping approach, based on Google News RSS feeds and the newspaper3k Python library², targeting news articles related to the 60 extreme weather events of interest. The scraper retrieved RSS feeds within a one-month time window before and after each event date, parsing content to extract relevant details such as article titles, descriptions, and full text, with a focus on English-language news articles. To isolate location-specific information, we first extracted sentences from the news articles and then employed the Flair NER tagger³ to identify and extract entities of the GPE (Geopolitical Entity) category. Sentences lacking location tags were considered noisy and subsequently filtered out, resulting in a curated dataset of 127,454 sentences. This process allowed us to focus on extreme weather-specific news, containing information about

¹<https://www.climameter.org>

²<https://newspaper.readthedocs.io/en/latest>

³<https://huggingface.co/flair/ner-english>

vulnerabilities, impacts, and emergency responses across diverse geographical locations. Table A.1 provides additional details on the selected events, including event types, dates, and associated locations.

We note that effectively retrieving news articles for a specific extreme weather event requires more than simply searching for a single attribute, such as the event name, since many articles either omit explicit mentions or discuss broader contexts. To address this issue, we constructed comprehensive search queries combining the event’s name or an appropriate proxy (e.g., “Typhoon Yagi” or “Central Asia heatwave”), the event location obtained from ClimaMeter, and the region of interest for the analysis, along with predefined thematic keywords categorized into: (a) *public*, (b) *economic*, and (c) *weather conditions*. Public keywords (e.g., “relief”, “evacuation”) emphasize community resilience and emergency responses, economic keywords (e.g., “economic cost”, “insurance”) target financial impacts and recovery, and weather conditions keywords (e.g., “high temperature”, “rain”) capture specific meteorological factors. A comprehensive list of adopted keywords is provided in Table A.2.

An example of the search query is: “*Typhoon Yagi Vietnam resilience weather*”, where we include the general keyword “*weather*” to improve relevance and context. Without this keyword, terms like “impact”, “resilience”, or “insurance” could yield results unrelated to weather, such as economic policies or general disaster recovery efforts. We also note that the time window adopted for our web scraping – i.e., ± 1 month with respect to a given extreme weather event – further anchors the search to relevant content specific to the event of interest.

3.2. Alignment Data

The dataset introduced in section 3.1 provides event-specific textual data, serving as a foundation for developing *ExtremeAlign*, a high-quality, task-specific alignment dataset leveraging LLMs. Training small language models to effectively handle diverse extreme weather analysis tasks necessitates extensive, accurately annotated data. However, manual annotation of large-scale datasets focused on extreme weather events is both labor-intensive and costly, making it impractical. To overcome this challenge, we utilize *Qwen2.5-32B-Instruct* [47], an expert-level model in extreme weather reasoning, to automatically generate high-quality alignment data through a one-shot prompting strategy, collectively denoted as $\mathcal{D}_{\text{one-shot}}$. This approach incorporates comprehensive definitions of task-specific categories and clear

guidelines within each prompt, enhancing data consistency, facilitating human alignment, and reducing subjective biases during dataset construction.

To effectively analyze extreme events, we design prompts for three interconnected tasks: *vulnerability/impact/emergency assessment*, *topic/subtopic labeling and keyword extraction*, and *emotion analysis*. The first assesses event severity and urgency, the second categorizes thematic content, and the third evaluates public sentiment and emotional responses. Together, these tasks provide a comprehensive framework capturing the vulnerability and impact, societal reactions, and broader implications of extreme weather.

Consequently, we formulate the overall dataset generation task as:

$$\mathcal{D}_{\text{gen}} \leftarrow LLM_p(T, \mathcal{D}_{\text{one-shot}})$$

where \mathcal{D}_{gen} represents the generated dataset, p indicates the prompt used during inference, and T refers to the three tasks.

Prior to prompt design, we considered several key aspects to ensure systematic categorization and consistency. Clear category definitions, including explicit inclusion and exclusion criteria, were established to avoid ambiguity. We also integrated structured reasoning into our prompts, requiring the model to articulate step-by-step thinking before assigning probabilities. To enforce interpretability and consistency, we constrained probability scores to sum exactly to one and adopted a strict output format with distinct `<think>` and `<output>` sections. A strict format standardizes output structure, which is essential for post-processing and evaluation. These considerations facilitated structured decision-making, improved response accuracy, and enhanced model reliability in real-world extreme weather analysis.

We followed a consistent prompting strategy across all three tasks:

Step 1: Given a sentence, we prompt the LLM to identify applicable categories based on the information taxonomy in Table B.3.

Step 2: The LLM assigns a probability score (between 0 and 1) to each category, reflecting its confidence.

Step 3: Probability scores across all categories must sum exactly to 1.

By framing these tasks as ranking problems, we enable the LLM to learn relative importance among categories, thereby reflecting the inherent ambiguity and complexity of extreme event-related information. Generating probability scores instead of single labels allows for representations of uncer-

tainty, accommodates overlapping categories, and supports accurate decision-making in critical applications such as disaster response, risk analysis, and resource allocation. Detailed prompts for each task are provided in section Appendix B.1.

The final dataset consists of 30,000 high-quality annotations, evenly divided across the three primary tasks, ensuring balanced supervision for multi-task alignment.

3.3. Extreme Weather Reasoning-Aware Alignment (EWRA)

To improve the performance of SLMs, we go beyond conventional supervised fine-tuning (SFT), which typically depends on human-annotated datasets. Instead, we adopt an advanced reasoning-based learning framework, a specialized variant of SFT, combined with a knowledge transfer mechanism from LLMs. This strategy involves training the SLM using a synthetic alignment dataset $\mathcal{D} = \{(p, q, r_p)\}$, generated by the LLM. Here, p denotes a prompt, q refers to an input query, and r_p is the corresponding output that incorporates reasoning, including both intermediate reasoning steps and the final task-specific output.

The output r_p is composed of two parts: the reasoning steps $r_{\text{reasoning}} = [t_1, t_2, \dots, t_k]$, where k is the total number of reasoning tokens, and the final output r_{final} , which is the ultimate answer or prediction. The reasoning process is generated step-by-step as $r_p = [r_{\text{reasoning}}, r_{\text{final}}]$, where $r_{\text{reasoning}}$ contains intermediate reasoning and r_{final} is the final output of the task.

At each time step j , the token t_j is sampled from the generation distribution $\pi_\phi(\cdot | s_j)$, where π_ϕ represents the model’s learned distribution over possible tokens, and s_j represents the model’s state. The training objective is to maximize the likelihood of the reasoning sequence conditioned on the prompt. We consider two distinct prompting regimes: 1. Explicit prompt incorporates detailed task definitions, inclusion/exclusion criteria, and step-by-step reasoning instructions. These promote structured alignment and precise instruction following. 2. Implicit prompt, in contrast, excludes such definitions. This setting encourages reliance on the LLM’s internalized domain knowledge and captures more flexible, context-dependent reasoning patterns.

The respective training objectives are defined as:

$$L_{\text{Explicit}}(\phi) = -\mathbb{E}_{(p,q,r_p) \sim \mathcal{D}_{\text{explicit}}} \left[\sum_{j=1}^k \log \pi_\phi(t_j | s_j, p) \right], \quad (1)$$

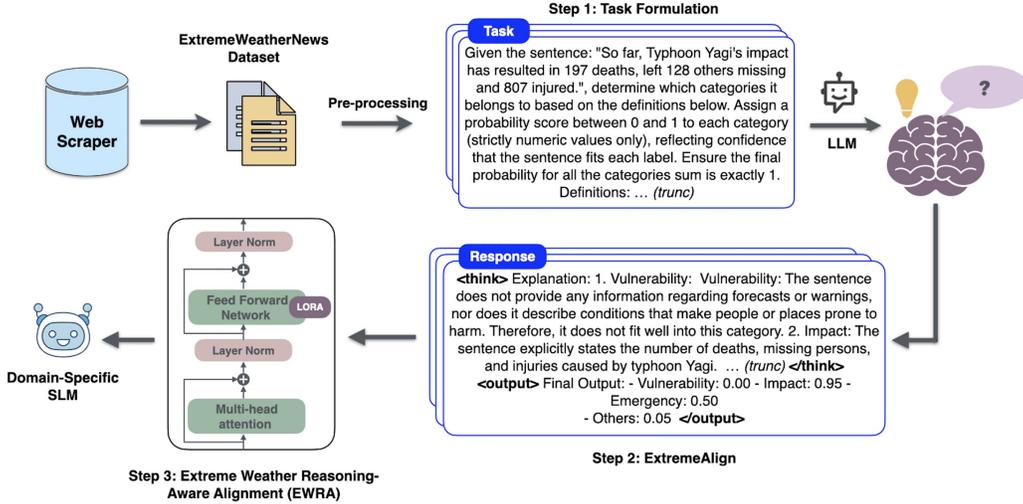


Figure 1: An overview of EWRA. The process begins with web scraping news articles to construct the ExtremeWeatherNews, followed by pre-processing. In Step 1, the task is formulated for an LLM. Step 2 introduces the ExtremeAlign (Alignment Data Construction), which structures reasoning-based outputs. Finally, in Step 3, a domain-specific Small Language Model (SLM) is trained using EWRA to improve reasoning alignment for extreme weather analysis.

$$L_{\text{Implicit}}(\phi) = -\mathbb{E}_{(p,q,r_p) \sim \mathcal{D}_{\text{implicit}}} \left[\sum_{j=1}^k \log \pi_{\phi}(t_j \mid s_j, p) \right]. \quad (2)$$

Our training strategy for *EWRA* follows a two-stage curriculum. We begin by training SLMs on data generated with implicit prompts. This stage promotes generalizable, context-aware reasoning by encouraging the model to draw upon latent knowledge from the LLM. We then fine-tune the model using explicitly guided examples that include structured definitions and intermediate reasoning steps. This phase sharpens the model’s ability to follow prompts and improves task-specific performance.

This curriculum style approach enables robust transfer of reasoning capabilities from LLMs to SLMs, blending the flexibility of implicit reasoning with the precision of explicit guidance. By leveraging both types of prompts, *EWRA* supports nuanced decision-making and alignment in high-stakes applications such as extreme weather analysis.

Overall, *EWRA* leverages reasoning-based alignment to enhance the reasoning capabilities of SLMs for extreme weather analysis. By utilizing *ExtremeAlign*, a structured alignment dataset generated using LLMs, we trans-

fer domain-specific reasoning properties to SLMs. In the following section, we present our experimental results, demonstrating the effectiveness of EWRA in improving reasoning performance for extreme weather event analysis tasks.

4. Experiments

4.1. Model Setup

We use Unsloth⁴ for training our models and use Unsloth’s pre-selected defaults for fine-tuning such as we set the LoRA rank to 16 and fine-tune all modules. The LoRA scaling factor is set to 16, with a dropout rate of 0 and no bias applied. Gradient checkpointing is enabled using “unsloth” for memory efficiency. The random seed is fixed at 3407 for reproducibility. Rank-stabilized LoRA is disabled, and loftq config is set to None. All the models are trained for one epoch with a learning rate of 2e-4 and a weight decay of 0.01. We employ a linear learning rate scheduler with 5 warmup steps. The effective batch size is 64, achieved by setting the per-device batch size to 16 and applying gradient accumulation over 4 steps. Training is optimized using AdamW 8-bit, with mixed precision enabled via FP16 or BF16, depending on hardware support. Max sequence length is set to 2048 with 4bit quantization to reduce memory usage. Learning rate scheduler type is linear. The experiments were conducted on a workstation equipped with four NVIDIA RTX 3090/4090 GPUs, each with 24GB of VRAM. For alignment data generation, Qwen2.5-32B-Instruct processes each sample in 1 minute, utilizing 21GB of VRAM. Task-specific models were trained with a batch size of 16, a maximum sequence length of 256, a learning rate of 2e-5, and a dropout rate of 0.3. For computing BERTScore, we use the default model employed *viz.* roberta-large. We adopt Query-Key (QK) fine-tuning, where only the query and key matrices in the attention layers are updated [48]. Our goal is to align SLMs to reason better about extreme weather events without overwriting their general language understanding. Since the value matrices and MLP layers are key to storing factual knowledge, freezing them preserves the LLM’s pretraining, while still allowing adaptation of attention dynamics. For our training protocol, we run each of the explicit and implicit alignment datasets for 2 epochs. For the EWRA strategy, we adopt a two-stage curriculum: the model is first trained for 1 epoch on the implicit data, followed by 1 epoch on the explicit data.

⁴<http://github.com/unslothai/unsloth>

4.2. Evaluation

Given the subjective nature of extreme weather analysis tasks, we use ranking-based evaluation instead of absolute accuracy to better capture model performance for our task, where multiple correct outputs exist. Specifically, we employ Spearman Rank Correlation [49] to assess alignment between model predictions and human annotations, ensuring relevance and expert-aligned ordering.

Spearman Rank Correlation. The Spearman rank correlation (SRC) coefficient (ρ) measures the strength and direction of the monotonic relationship between predicted and ground-truth rankings. For two ranked sets X and Y with n observations, ρ is computed as:

$$\rho = 1 - \frac{6 \sum_{i=1}^n d_i^2}{n(n^2 - 1)}$$

where d_i represents the rank difference between the i -th observation in X and Y . Spearman’s ρ ranges from -1 (perfect negative correlation) to +1 (perfect positive correlation), with 0 indicating no monotonic association.

Similarly, to evaluate the generated explanations, we use the Jaccard Index, which quantifies the overlap between two sets of words or tokens representing the explanations. It evaluates how similar the content of two explanations is by comparing their shared elements. It is calculated as the ratio of the number of common words (or tokens) between two explanations to the total number of unique words across both explanations. We also report BERTScore[50], which is a token-level similarity metric that evaluates the quality of generated text by computing contextualized embeddings using a pre-trained transformer model. Instead of relying on exact word matching, BERTScore measures semantic similarity between candidate and reference texts by computing cosine similarity between their embeddings. The range for both metrics is from 0 to 1.

4.3. Model Description

We compare the performance of different models on the *ExtremeWeatherNews* dataset across three key tasks: (1) Vulnerability/Impact/Emergency Statement Assessment, (2) Topic/Subtopic Labeling and Keyword Extraction, (3) Emotion Analysis. These tasks are evaluated using the alignment data described in section 3.2.

Below, we describe the models used in our experiments.

1. **Task-Specific Models:** For the first task, we use a RoBERTa-based model [51] and perform multi-class classification. As this task has not been previously explored, we utilize the model as a baseline for our experiments. For emotion analysis, we leverage the proposed model by [52] for our experiments. We also record the probabilities over all the classes. Similarly, for the third task, we also employ RoBERTa, since this task has not been previously formulated as a ranking problem, while for keyword extraction, we utilize concept parsing [53, 54] which utilizes part-of-speech (POS) tagging and dependency parsing to identify syntactic relationships. This model is a fine-tuned transformer-based classifier optimized individually for each task. It does not leverage chain-of-thought explanations but is trained specifically for classification using supervised learning.
2. **Qwen2.5 Instruct Models (1.5B, 3B):** Qwen2.5 [47] is a family of LLMs developed by Alibaba’s Qwen team. We experiment with two variants: Qwen2.5-1.5B-Instruct, and Qwen2.5-3B-Instruct. These models differ in capacity, with larger variants expected to capture more complex reasoning patterns and provide more reliable explanations.
 - (a) **Zero-shot setting:** We established a baseline by performing zero-shot inference on the ExtremeAlign dataset using the Qwen2.5-1.5B-Instruct model without any fine-tuning. This provides a reference point to assess the effectiveness of our alignment strategies.
 - (b) **Direct-SFT:** A standard supervised fine-tuning (SFT) approach where the model is trained only on the final task-specific outputs (category labels). While category definitions are included in the instruction, the training data omits any reasoning traces or intermediate steps.
 - (c) **ReasonImplicit-SFT:** In this setting, the model is trained on both reasoning traces and final category labels. However, task-specific definitions are excluded from the instructions. This design encourages the model to rely on its internal parametric knowledge and domain understanding during reasoning.
 - (d) **ReasonExplicit-SFT:** The model is trained using both step-by-step reasoning traces and final outputs, with the instructions explicitly prompting for structured reasoning. Additionally, detailed task

definitions are included in the prompt to guide the model’s predictions.

- (e) EWRA (Extreme Weather Reasoning-Aware Alignment): Our proposed alignment strategy described in section 3.3. It employs a two-stage curriculum. First, the model is trained on implicitly prompted data (without definitions) to encourage generalizable and context-aware reasoning. This is followed by fine-tuning on explicitly guided data (with structured definitions and reasoning), promoting task adherence and instruction following.

4.4. Data Splits and Gold TestSet

The ExtremeAlign dataset is partitioned into training, validation, and test splits using a 70%, 15%, and 15% ratio, respectively. To evaluate model performance under realistic human-like reasoning, we additionally construct a gold-standard test set comprising 600 samples manually annotated by human experts, 200 for each core task: vulnerability/impact/emergency assessment, topic/subtopic/keyword labeling, and emotion analysis. The annotation process followed a detailed set of guidelines, as described in section Appendix B.2. These annotations reflect both explicit instruction adherence and implicit domain knowledge, capturing the nuanced decision-making typically employed by humans in extreme weather analysis.

4.5. Results & Discussion

Tables 1, 2 and 3 report the performance of various models on the ExtremeAlign dataset across three tasks. Evaluation is conducted using Spearman’s rank correlation coefficient (ρ) to assess ranking accuracy, while explanation quality is measured using the Jaccard Index and BERTScore by comparing model outputs with the gold set.

Overall, EWRA consistently outperforms other training strategies, particularly when applied to the Qwen2.5-3B-Instruct model with higher number of parameters. Compared to ReasonExplicit-SFT, EWRA provides clear improvements across most tasks and metrics. For instance, in the Vulnerability/Impact/Emergency assessment task on the 3B model, EWRA improves Jaccard Index by 2.7% and increases SRC score by 4.3%. These gains indicate that reasoning alignment helps models generalize better and make more precise predictions. EWRA and ReasonExplicit-SFT are the top-performing fine-tuning strategies overall. While improvements in generated

Table 1: Performance comparison across Vulnerability/Impact/Emergency Assessment task using different fine-tuning strategies on Qwen2.5-1.5B and Qwen2.5-3B models. Metrics include Spearman Rank Correlation (SRC (ρ)), Jaccard Index, and BERTScore.

Vulnerability/Impact/Emergency Assessment				
Model	Method	SRC (ρ)	Jaccard Index	BERTScore
Task-Specific Models	-	0.1467	-	-
Qwen2.5-1.5B-Instruct	Zero-shot	0.3638	0.1277	0.7133
	Direct-SFT	0.7290	-	-
	ReasonImplicit-SFT	0.4200	0.1861	0.7290
	ReasonExplicit-SFT	0.7960	0.2585	0.8897
	EWRA	0.7140	0.2573	0.9028
Qwen2.5-3B-Instruct	Zero-shot	0.5400	0.1226	0.3744
	Direct-SFT	0.7910	-	-
	ReasonImplicit-SFT	0.4860	0.1937	0.6993
	ReasonExplicit-SFT	0.7820	0.3294	0.9163
	EWRA	0.8230	0.3340	0.9161

explanation (measured via BERTScore) are less significant as compared to ReasonExplicit-SFT, this is likely due to the fact that these metrics emphasize semantic overlap, which may remain high even when explanation structure or specificity differs. As such, BERTScore may under-represent the qualitative differences in reasoning alignment captured by EWRA.

1. Zero-shot vs EWRA: zero-shot models typically underperform compared to EWRA and other fine-tuned variants—due to their lack of task-specific alignment and limited domain grounding. However, we may note that Qwen2.5-1.5B-Instruct in the zero-shot setting achieved a Spearman correlation of 0.4104/0.6702, while the EWRA-aligned model reached only 0.5673/0.6671. A similar trend is observed for Qwen2.5-3B-Instruct, where the zero-shot model attained 0.4660/0.7727, outperforming the EWRA variant (0.5386/0.7149). While EWRA is effective for broad topic labeling and structured categorization, it falls short in subtopic labeling, which often requires nuanced semantic understanding and sensitivity to fine-grained distinctions. One likely reason is the task design: subtopic evaluation relies on ranking-based accuracy metrics, which are sensitive to the model’s ability to produce all relevant labels. However, SFT-aligned models—particularly under EWRA—tend to predict only a subset of categories, reducing performance on ranking-based evaluations. This is further supported by explanation similarity metrics (e.g., BERTScore), where ReasonEx-

Table 2: Performance comparison across Topic/Subtopic Labeling and Keyword Extraction task, using different fine-tuning strategies on Qwen2.5-1.5B and Qwen2.5-3B models. Metrics include Spearman Rank Correlation (SRC (ρ)), Jaccard Index, and BERTScore.

Topic/Subtopic Labeling				
Model	Method	SRC (ρ)	Jaccard Index	BERTScore
Task-Specific Models	-	0.2100/0.1999	-	-
Qwen2.5-1.5B-Instruct	Zero-shot	0.4104/ 0.6702	0.0436/0.0436	0.6342/0.6342
	Direct-SFT	0.3264/0.5770	-/-	-/-
	ReasonImplicit-SFT	0.3040/0.52883	0.1670/0.1670	0.8741/0.8741
	ReasonExplicit-SFT	0.5157/0.6589	0.1744/0.1720	0.8772/0.8772
	EWRA	0.5673/0.6671	0.1711/0.1711	0.8765/0.8765
Qwen2.5-3B-Instruct	Zero-shot	0.4660/ 0.7727	0.1592/0.1592	0.8257/0.8257
	Direct-SFT	0.4179/0.6397	-/-	-/-
	ReasonImplicit-SFT:	0.4131/0.6066	0.1907/0.1671	0.8831/0.8741
	ReasonExplicit-SFT	0.5692/0.6527	0.2045/0.1743	0.8879/0.8771
	EWRA	0.5386/0.7149	0.2026/ 0.2026	0.8890/0.8890
Keyword Extraction				
Model	Method	Jaccard Index		
Task-Specific Models	-	0.132		
Qwen2.5-1.5B-Instruct	Zero-shot	0.1257		
	Direct-SFT	0.1984		
	ReasonImplicit-SFT	0.2813		
	ReasonExplicit-SFT	0.2343		
	EWRA	0.2333		
Qwen2.5-3B-Instruct	Zero-shot	0.2817		
	Direct-SFT	0.2698		
	ReasonImplicit-SFT:	0.3269		
	ReasonExplicit-SFT	0.2343		
	EWRA	0.2856		

licit and EWRA settings achieve the highest scores, indicating that while reasoning quality improves, the alignment objective may suppress broader coverage needed for exhaustive subtopic labeling.

- ReasonImplicit-SFT vs ReasonExplicit-SFT: A comparison between ReasonImplicit-SFT and ReasonExplicit-SFT highlights the importance of including explicit definitions and step-by-step guidance. While gains in Emotion analysis are relatively small (e.g., BERTScore improvement of 0.15%), Vulnerability/Impact/Emergency assessment task shows a 70.3% relative improvement in Jaccard when moving from implicit to explicit reasoning on the 3B model. This suggests that when categories are fine-grained or semantically overlapping, models benefit from being explicitly told how to structure their reasoning. However, we observe

Table 3: Performance comparison across Emotion Analysis task using different fine-tuning strategies on Qwen2.5-1.5B and Qwen2.5-3B models. Metrics include Spearman Rank Correlation (SRC (ρ)), Jaccard Index, and BERTScore.

Emotion Analysis				
Model	Method	SRC (ρ)	Jaccard Index	BERTScore
Task-Specific Models	-	0.2876	-	-
Qwen2.5-1.5B-Instruct	Zero-shot	0.6071	0.1336	0.7434
	Direct-SFT	0.8389	-	-
	ReasonImplicit-SFT:	0.7828	0.2528	0.8986
	ReasonExplicit-SFT	0.5578	0.2583	0.8979
	EWRA	0.8308	0.2565	0.8976
Qwen2.5-3B-Instruct	Zero-shot	0.6458	0.1112	0.2666
	Direct-SFT	0.8442	-	-
	ReasonImplicit-SFT:	0.8380	0.2627	0.9011
	ReasonExplicit-SFT	0.8716	0.2646	0.8989
	EWRA	0.8708	0.2664	0.8995

that results on keyword extraction reveal that ReasonImplicit-SFT outperforms ReasonExplicit-SFT in several cases, especially on the smaller Qwen2.5-1.5B model. This can be attributed to the flexible nature of the task: since keyword extraction does not rely on fixed category definitions, models may benefit from a more latent understanding of task intent, which ReasonImplicit-SFT encourages. The implicit formulation likely gives the model room to adapt and use contextual cues from the input without being overly constrained by rigid definitions. This suggests that ReasonImplicit-SFT may promote better domain adaptation in tasks like keyword extraction, where the model must infer task structure from context rather than from explicit templates.

- ReasonExplicit-SFT vs Direct SFT: Direct-SFT achieves a high Spearman correlation (e.g., 0.8389 for the 3B model) on the emotion analysis task, likely due to the dominance of a few frequent classes such as sadness, and fear. While this indicates some alignment with ground-truth rankings, the absence of reasoning traces suggests that the model may be overfitting to surface-level patterns or task-specific heuristics rather than learning meaningful structure. In contrast, EWRA and ReasonExplicit-SFT consistently performs well across all metrics, demonstrating its ability to generate outputs that are not only accurate but also coherent and well-aligned with the reasoning behind the task.

4.6. Effectiveness of EWRA

To further examine the role of reasoning alignment in model performance, we perform an ablation study by introducing a ReverseEWRA (Table 4) setting. ReverseEWRA is a contrastive variant of EWRA in which the curriculum order is reversed: the model is first trained on explicit instructions and reasoning, followed by fine-tuning using the implicit setting. This baseline helps isolate the impact of curriculum ordering on model alignment and reasoning capabilities. The motivation behind this experiment is to assess whether the performance gains in EWRA stem from the curriculum structure rather than just prolonged fine-tuning.

Table 4: Effectiveness of EWRA: Task-wise Ablation using ReverseEWRA

Vulnerability/Impact/Emergency Assessment				
Model	Method	SRC (ρ)	Jaccard Index	BERTScore
Qwen2.5-1.5B-Instruct	ReverseEWRA	0.7080	0.2555	0.9025
Qwen2.5-3B-Instruct	ReverseEWRA	0.7070	0.3117	0.9124
Topic/Subtopic Labeling				
Model	Method	SRC (ρ)	Jaccard Index	BERTScore
Qwen2.5-1.5B-Instruct	ReverseEWRA	0.5182 / 0.4949	0.1595 / 0.1595	0.8729 / 0.8728
Qwen2.5-3B-Instruct	ReverseEWRA	0.5209 / 0.6747	0.1931 / 0.1931	0.8859 / 0.8859
Keyword Extraction				
Model	Method	Jaccard Index		
Qwen2.5-1.5B-Instruct	ReverseEWRA	0.2411		
Qwen2.5-3B-Instruct	ReverseEWRA	0.2939		
Emotion Analysis				
Model	Method	SRC (ρ)	Jaccard Index	BERTScore
Qwen2.5-1.5B-Instruct	ReverseEWRA	0.8650	0.2512	0.8974
Qwen2.5-3B-Instruct	ReverseEWRA	0.8543	0.2667	0.9005

Across all tasks, particularly those involving multi-label ranking, such as Vulnerability/Impact/Emergency assessment and Topic/Subtopic labeling, we observe notable performance degradation in the ReverseEWRA setting. For example, on the Qwen2.5-3B model, the SRC score for Vulnerability/Impact/Emergency assessment drops from 0.8230 in EWRA to 0.7070 in ReverseEWRA. ReverseEWRA leads to a consistent drop in SRC for Topic/Subtopic labeling across both models, with a relative decline of approximately 8.7% (from 0.5673 to 0.5182) and 25.8% (from 0.6671 to 0.4949) for Qwen2.5-1.5B, and 3.3% (from 0.5386 to 0.5209) and 5.6% (from 0.7149 to 0.6747) for Qwen2.5-3B.

Interestingly, the Qwen2.5-1.5B model achieves its best performance in Emotion Analysis under the ReverseEWRA setting, with an SRC score of 0.8650—outperforming all other fine-tuning strategies for that task. This suggests that in some domains (like emotion inference), starting with explicit reasoning may provide a better inductive bias for low-capacity models, even if the overall curriculum is less optimal across tasks.

5. Conclusion

In this study, we presented Extreme Weather Reasoning-Aware Alignment (EWRA), a novel fine-tuning framework that transfers the structured reasoning capabilities of LLMs to small language models (SLMs) for the domain of extreme weather event analysis. By integrating LLM-generated reasoning paths into a two-stage training strategy, EWRA enables SLMs to develop both context-aware and task-specific reasoning skills. We introduced two new datasets: ExtremeWeatherNews, a rich collection of extreme weather event news data, and ExtremeAlign, a reasoning-augmented alignment set tailored to three key analysis tasks: vulnerability/impact/emergency assessment, topic/subtopic/keyword labeling, and emotion analysis. Our experiments across these tasks demonstrate that EWRA substantially improves the alignment of SLMs with domain-specific reasoning, outperforming both standard supervised fine-tuning and other reasoning-based approaches, particularly on models such as Qwen2.5-3B-Instruct with higher number of parameters as compared to Qwen2.5-1.5B-Instruct model. As future work, we plan to extend EWRA to broader generative tasks such as temporal summarization and scenario simulation, while also investigating cross-task generalization by developing modular, instruction-tuned variants of ExtremeAlign.

References

- [1] V. Masson-Delmotte, P. Zhai, S. Pirani, C. Connors, S. Péan, N. Berger, Y. Caud, L. Chen, M. Goldfarb, P. M. Scheel Monteiro, IPCC, 2021: Summary for Policymakers. In: *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change (2021)*.
- [2] R. Mao, Q. Lin, Q. Liu, G. Mengaldo, E. Cambria, Understanding public perception towards weather disasters through the lens of metaphor,

- in: Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, (IJCAI-24), 2024, pp. 7394–7402.
- [3] G. Camps-Valls, M.-Á. Fernández-Torres, K.-H. Cohrs, A. Höhl, A. Castelletti, A. Pacal, C. Robin, F. Martinuzzi, I. Papoutsis, I. Prapas, et al., Artificial intelligence for modeling and understanding extreme weather and climate events, *Nature Communications* 16 (1) (2025) 1919.
 - [4] M. J. Hammond, A. S. Chen, S. Djordjević, D. Butler, O. Mark, Urban flood impact assessment: A state-of-the-art review, *Urban Water Journal* 12 (1) (2015) 14–29.
 - [5] F. Zennaro, E. Furlan, C. Simeoni, S. Torresan, S. Aslan, A. Critto, A. Marcomini, Exploring machine learning potential for climate change risk assessment, *Earth-Science Reviews* 220 (2021) 103752.
 - [6] A. Coletti, P. D. Howe, B. Yarnal, N. J. Wood, A support system for assessing local vulnerability to weather and climate, *Natural hazards* 65 (2013) 999–1008.
 - [7] M. M. Bechtel, M. Mannino, Ready when the big one comes? natural disasters and mass support for preparedness investment, *Political behavior* 45 (3) (2023) 1045–1070.
 - [8] S. H. Eriksen, P. M. Kelly, Developing credible vulnerability indicators for climate adaptation policy assessment, *Mitigation and adaptation strategies for global change* 12 (2007) 495–524.
 - [9] K. Ong, R. Mao, D. Varshney, R. Satapathy, E. Cambria, G. Mengaldo, Sentiment analysis on climate change for sustainable investment, in: Proceedings of ICDM Workshops, 2024, pp. 479–486.
 - [10] N. Li, S. Zahra, M. Brito, C. Flynn, O. Görnerup, K. Worou, M. Kurfali, C. Meng, W. Thiery, J. Zscheischler, et al., Using llms to build a database of climate extreme impacts, in: Proceedings of the 1st Workshop on Natural Language Processing Meets Climate Change (ClimateNLP 2024), 2024, pp. 93–110.
 - [11] J. Bulian, M. S. Schäfer, A. Amini, H. Lam, M. Ciaramita, B. Gaiarin, M. C. Huebscher, C. Buck, N. G. Mede, M. Leippold, et al., Assessing

- large language models on climate information, in: International Conference on Machine Learning, PMLR, 2024, pp. 4884–4935.
- [12] Y. Wang, H. A. Karimi, Exploring large language models for climate forecasting, arXiv preprint arXiv:2411.13724 (2024).
 - [13] H. Zhu, P. Tiwari, Climate change from large language models, arXiv preprint arXiv:2312.11985 (2023).
 - [14] D. Zheng, M. Lapata, J. Z. Pan, Large language models as reliable knowledge bases?, arXiv preprint arXiv:2407.13578 (2024).
 - [15] D. Thulke, Y. Gao, P. Pelsler, R. Brune, R. Jalota, F. Fok, M. Ramos, I. van Wyk, A. Nasir, H. Goldstein, et al., Climategpt: Towards ai synthesizing interdisciplinary research on climate change. arxiv. 2024 doi: 10.48550, arXiv preprint arXiv.2401.09646.
 - [16] S. A. Vaghefi, D. Stambach, V. Muccione, J. Bingler, J. Ni, M. Kraus, S. Allen, C. Colesanti-Senni, T. Wekhof, T. Schimanski, et al., Chatclimate: Grounding conversational ai in climate science, *Communications Earth & Environment* 4 (1) (2023) 480.
 - [17] H. Turbé, M. Bjelogrić, C. Lovis, G. Mengaldo, Evaluation of post-hoc interpretability methods in time-series classification, *Nature Machine Intelligence* 5 (3) (2023) 250–260.
 - [18] G. Mengaldo, Progress and challenges in the study of extreme weather, EGU25 (EGU25-7685) (2025).
 - [19] M. M. de Brito, J. Sodoge, A. Fekete, M. Hagenlocher, E. Koks, C. Kuhlicke, G. Messori, M. de Ruiter, P.-J. Schweizer, P. J. Ward, Uncovering the dynamics of multi-sector impacts of hydrological extremes: A methods overview, *Earth’s Future* 12 (1) (2024) e2023EF003906.
 - [20] J. Sodoge, C. Kuhlicke, M. D. Mahecha, M. M. de Brito, Text mining uncovers the unique dynamics of socio-economic impacts of the 2018–2022 multi-year drought in germany, *Natural Hazards and Earth System Sciences* 24 (5) (2024) 1757–1777.

- [21] P. H. Alencar, J. Sodoge, E. Paton, M. Madruga de Brito, Flash droughts and their impacts—using newspaper articles to assess the perceived consequences of rapidly emerging droughts, *Environmental Research Letters* (2024).
- [22] B. Zhang, F. Schilder, K. H. Smith, M. J. Hayes, S. Harms, T. Tadesse, Tweetdrought: A deep-learning drought impacts recognizer based on twitter data, *arXiv preprint arXiv:2212.04001* (2022).
- [23] M. Kahle, M. Kempf, B. Martin, R. Glaser, Classifying the 2021 ‘ahrta’ flood event using hermeneutic interpretation, natural language processing, and instrumental data analyses, *Environmental Research Communications* 4 (5) (2022) 051002.
- [24] N. Webersinke, M. Kraus, J. A. Bingler, M. Leippold, Climatebert: A pretrained language model for climate-related text, *arXiv preprint arXiv:2110.12010* (2021).
- [25] E. Cambria, X. Zhang, R. Mao, M. Chen, K. Kwok, SenticNet 8: Fusing emotion AI and commonsense AI for interpretable, trustworthy, and explainable affective computing, in: *Proceedings of International Conference on Human-Computer Interaction (HCII)*, Washington DC, USA, 2024, pp. 197–216.
- [26] E. Cambria, R. Mao, M. Chen, Z. Wang, S.-B. Ho, Seven pillars for the future of artificial intelligence, *IEEE Intelligent Systems* 38 (6) (2023) 62–69.
- [27] C. Duong, V. C. Raghuram, A. Lee, R. Mao, G. Mengaldo, E. Cambria, Neurosymbolic ai for mining public opinions about wildfires, *Cognitive Computation* 16 (4) (2024) 1531–1553.
- [28] F. S. Varini, J. Boyd-Graber, M. Ciaramita, M. Leippold, Climatext: A dataset for climate change topic detection, *arXiv preprint arXiv:2012.00483* (2020).
- [29] D. Chen, Y. Zheng, L. Ma, F. Qin, An ontology-based knowledge representation method for typhoon events from chinese news reports, *ISPRS International Journal of Geo-Information* 12 (9) (2023) 383.

- [30] R. Mao, X. Li, K. He, M. Ge, E. Cambria, MetaPro Online: A computational metaphor processing online system, in: Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations), 2023, pp. 127–135.
- [31] C. Zhou, P. Liu, P. Xu, S. Iyer, J. Sun, Y. Mao, X. Ma, A. Efrat, P. Yu, L. Yu, et al., Lima: Less is more for alignment, *Advances in Neural Information Processing Systems* 36 (2023) 55006–55021.
- [32] G. Colverd, P. Darm, L. Silverberg, N. Kasmanoff, Floodbrain: Flood disaster reporting by web-based retrieval augmented generation with an llm, *arXiv preprint arXiv:2311.02597* (2023).
- [33] H. T. Otal, E. Stern, M. A. Canbaz, Llm-assisted crisis management: Building advanced llm platforms for effective emergency response and public collaboration, in: 2024 IEEE Conference on Artificial Intelligence (CAI), IEEE, 2024, pp. 851–859.
- [34] K. Yin, C. Liu, A. Mostafavi, X. Hu, Crisissense-llm: Instruction fine-tuned large language model for multi-label social media text classification in disaster informatics, *arXiv preprint arXiv:2406.15477* (2024).
- [35] W. Chen, J. Fang, Optimizing ai-driven disaster management through llms (2024).
- [36] V. G. Goecks, N. R. Waytowich, Disasterresponsegpt: Large language models for accelerated plan of action development in disaster response scenarios, *arXiv preprint arXiv:2306.17271* (2023).
- [37] S. A. Vaghefi, C. Huggel, V. Muccione, H. Khashehchi, M. Leippold, Deep climate change: A dataset and adaptive domain pre-trained language models for climate change related tasks, in: *NeurIPS 2022 workshop on tackling climate change with machine learning*, 2022.
- [38] S. Mullappilly, A. Shaker, O. Thawakar, H. Cholakkal, R. Anwer, S. Khan, F. Khan, Arabic mini-climategpt: A climate change and sustainability tailored arabic llm, in: *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2023, pp. 14126–14136.
- [39] L. Ranaldi, A. Freitas, Aligning large and small language models via chain-of-thought reasoning, in: *Proceedings of the 18th Conference of*

the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers), 2024, pp. 1812–1827.

- [40] H. Wu, Y. Zhang, Z. Han, Y. Hou, L. Wang, S. Liu, Q. Gong, Y. Ge, Cot-driven framework for short text classification: Enhancing and transferring capabilities from large to smaller model, *Knowledge-Based Systems* (2025) 113057.
- [41] N. Ho, L. Schmid, S.-Y. Yun, Large language models are reasoning teachers, *arXiv preprint arXiv:2212.10071* (2022).
- [42] G. Hinton, O. Vinyals, J. Dean, Distilling the knowledge in a neural network, in: *Proceedings of the Neural Information Processing Systems (NeurIPS)*, Vol. 28, 2015, pp. 1–9.
URL <https://arxiv.org/abs/1503.02531>
- [43] V. Sanh, L. Debut, J. Chaumond, T. Wolf, Distilbert, a distilled version of bert: Smaller, faster, cheaper and lighter, in: *Proceedings of the 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing*, 2019.
URL <https://arxiv.org/abs/1910.01108>
- [44] R. Agarwal, N. Vieillard, Y. Zhou, P. Stanczyk, S. R. Garea, M. Geist, O. Bachem, On-policy distillation of language models: Learning from self-generated mistakes, in: *The Twelfth International Conference on Learning Representations*, 2024.
- [45] Y. Gao, Y. Xiong, X. Gao, K. Jia, J. Pan, Y. Bi, Y. Dai, J. Sun, H. Wang, H. Wang, Retrieval-augmented generation for large language models: A survey, *arXiv preprint arXiv:2312.10997 2* (2023).
- [46] D. Faranda, G. Messori, E. Coppola, T. Alberti, M. Vrac, F. Pons, P. Yiou, M. Saint Lu, A. N. Hisi, P. Brockmann, et al., Climameter: contextualizing extreme weather in a changing climate, *Weather and Climate Dynamics* 5 (3) (2024) 959–983.
- [47] A. Yang, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Li, D. Liu, F. Huang, H. Wei, H. Lin, J. Yang, J. Tu, J. Zhang, J. Yang, J. Yang, J. Zhou, J. Lin, K. Dang, K. Lu, K. Bao, K. Yang, L. Yu, M. Li, M. Xue, P. Zhang, Q. Zhu, R. Men, R. Lin, T. Li, T. Xia, X. Ren, X. Ren, Y. Fan,

- Y. Su, Y. Zhang, Y. Wan, Y. Liu, Z. Cui, Z. Zhang, Z. Qiu, Qwen2.5 technical report, arXiv preprint arXiv:2412.15115 (2024).
- [48] S. Goyal, C. Baek, J. Z. Kolter, A. Raghunathan, Context-parametric inversion: Why instruction finetuning can worsen context reliance, in: ICLR 2025 Workshop on Navigating and Addressing Data Problems for Foundation Models.
- [49] J. H. Zar, Spearman rank correlation, Encyclopedia of biostatistics 7 (2005).
- [50] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, Y. Artzi, Bertscore: Evaluating text generation with bert, arXiv preprint arXiv:1904.09675 (2019).
- [51] Y. Liu, Roberta: A robustly optimized bert pretraining approach, arXiv preprint arXiv:1907.11692 364 (2019).
- [52] A. Arora, E. Turcan, Evaluating the effectiveness of data augmentation for emotion classification in low-resource settings, arXiv preprint arXiv:2406.05190 (2024).
- [53] K. Ong, R. Mao, F. Xing, R. Satapathy, J. Sulaeman, E. Cambria, G. Mengaldo, ESGSenticNet: A neurosymbolic knowledge base for corporate sustainability analysis, arXiv preprint arXiv:2501.15720 (2025).
- [54] E. Cambria, R. Mao, S. Han, Q. Liu, Sentic parser: A graph-based approach to concept extraction for sentiment analysis, in: 2022 IEEE International Conference on Data Mining Workshops (ICDMW), IEEE, 2022, pp. 413–420.
- [55] M. S. Islam, A. Proma, Y. Zhou, S. N. Akter, C. Wohn, E. Hoque, Knowenvironment: an automated knowledge graph for climate change and environmental issues, in: AAAI 2022 Fall Symposium: the Role of AI in Responding to Climate Challenges, 2022.
- [56] P. Mishra, R. Mittal, Neuralnere: Neural named entity relationship extraction for end-to-end climate change knowledge graph construction, in: Tackling climate change with machine learning workshop at ICML, 2021.

Appendix A. ExtremeWeatherNews Dataset

This appendix section describes the ExtremeWeatherNews dataset and the steps used for data collection. The dataset comprises a curated corpus of news articles and reports centered around major extreme weather incidents. Events were selected using Climameter’s event lists [46], ensuring broad geographic and thematic coverage. Table A.1 lists the selected events used for news collection and additional details on the selected events, including event types, dates, and associated locations. To streamline data scraping, we implemented a multithreaded script capable of handling a high volume of concurrent requests. For targeted retrieval of relevant content, we developed a carefully curated set of keywords (see Table A.2). These keywords played a critical role in capturing a comprehensive and diverse range of articles. To ensure geographical relevance, we applied a filtering process based on location mentions. Specifically, we used administrative codes from the GeoNames⁵ database—namely admin1code (primary divisions such as states or provinces) and admin2code (secondary divisions like cities or districts). We retained only those articles whose mentioned locations matched the geographical scope of the corresponding events. Articles were further filtered to exclude sentences shorter than 30 characters or longer than 200 characters. Finally, to reduce noise and eliminate content unrelated to extreme weather, human annotators manually reviewed the data. They removed irrelevant samples, particularly those referring to unrelated domains. Further details about the annotation process are provided in Section Appendix B.2.

Appendix B. Methodology

We begin by developing a robust taxonomy, establishing a foundational structure that organizes climate concepts into key topics, subtopics, and specific entities. This taxonomy covers areas such as vulnerabilities, impacts, and emergency responses, creating a data-driven framework that supports integration with broader climate analysis tools. By categorizing data across these domains, the taxonomy enables dynamic, multi-dimensional exploration of insights, ranging from localized effects to wider climate patterns. However, empirical approaches in the literature are limited [55, 56], and these methods do not involve human evaluation to ensure the meaningfulness of

⁵<https://www.geonames.org>

Table A.1: Summary of Events with Types, Locations, and Dates.

Event Name	Event Type	Country	Event Date
Romania Floods	floods	Romania	31/08/2024
Poland Floods	floods	Poland	18/08/2024
USA Winter Storm	cold	USA	14/01/2024
Scandinavian Cold Spell	cold	Norway	08/01/2024
Ciro Snowstorm	cold	Germany	02/12/2023
North American Winter Storm	cold	Canada	27/12/2022
Hurricane Helene	wind	USA	27/09/2024
Typhoon Yagi	wind	Vietnam	08/09/2024
Storm Ingunn	wind	Norway	01/02/2024
Cyclone Belal	wind	France	15/01/2024
Cyclone Jasper	wind	Australia	18/12/2023
Storm Ciaran	wind	France	03/11/2023
Cyclone Tej	wind	Yemen	23/10/2023
Storms Babet and Aline	wind	Scotland	20/10/2023
Storm Poly	wind	Denmark	05/07/2023
Hurricane Ian Landfall	wind	USA	28/09/2022
April 2020 USA Tornado Outbreak	wind	USA	12/04/2020
Hurricane Irma Caribbean Landfall	wind	Caribbean	07/09/2017
European Heatwave	heatwave	Albania	19/07/2024
Eastern United States Heatwave	heatwave	USA	23/06/2024
Saudi Arabia Heatwave	heatwave	Saudi Arabia	18/06/2024
Eastern Mediterranean Heatwave	heatwave	Cyprus	14/06/2024
India Heatwave	heatwave	India	29/05/2024
Easter Extreme Weather in Europe	heatwave	Italy	01/04/2024
Morocco Heatwave	heatwave	Morocco	15/02/2024
Central Asia Heatwave	heatwave	Pakistan	30/11/2023
Brazil Heatwave	heatwave	Brazil	19/11/2023
October Heatwave in Europe	heatwave	Switzerland	13/10/2023
September Heatwave in Southern and Central Europe	heatwave	France	10/09/2023
Late Summer French Heatwave	heatwave	France	23/08/2023
Western USA Heatwave	heatwave	USA	31/07/2023
Cerberus Heatwave in Southern Europe	heatwave	Italy	25/07/2023
Southeast Asia Heat Peak	heatwave	Thailand	15/04/2023
Italy Multiple Floods	rain	Italy	19/10/2024
Storm Kirk	rain	France	09/10/2024
Storm Boris	rain	Austria	15/09/2024
Hurricane Beryl	rain	Jamaica	03/07/2024
Genoa Low Summer Floods	rain	France	24/06/2024
Bavaria Floods	rain	Germany	03/06/2024
Texas Floods	rain	USA	05/05/2024
South Brazil Floods	rain	Brazil	02/05/2024
China Floods	rain	China	23/04/2024
Dubai Floods	rain	UAE	16/04/2024
Storm Monica	rain	France	09/03/2024
California Floods	rain	USA	01/02/2024
San Diego Floods	rain	USA	01/22/2024
North-West USA and Canada Atmospheric River	rain	USA	06/12/2023
France and Italy Floods	rain	France	21/11/2023
Hurricane Otis	rain	Mexico	25/10/2023
New York Floods	rain	USA	29/09/2023
Mediterranean Depression Elias	rain	Greece	27/09/2023
Cape Town Floods	rain	South Africa	25/09/2023
Cevennes Floods	rain	France	17/09/2023
Medicane Daniel	rain	Lybia	11/09/2023
Guangdong and Hong Kong Floods	rain	Hong Kong	08/09/2023
Mediterranean Depression Daniel	rain	Greece	05/09/2023
Mediterranean Depression Rea	rain	Italy	29/08/2023
Storm Hans in Scandinavia	rain	Norway	08/08/2023
California Atmospheric River	rain	USA	10/01/2023
Medicane Ianos	rain	Greece	18/09/2020

constructed climate-related topics. Consequently, BERTopic is run on cli-

Table A.2: Keywords for Extreme Weather Event Search

Public Keywords	Economic Keywords	Weather Conditions Keywords
public, community, people, infrastructure, society, impact, disruption, affected, resilience, support, relief, aid, assistance, emergency, response, preparedness, adaptation, mitigation, awareness, engagement, cooperation, solidarity, social, health, welfare, equity, inclusion, vulnerability, risk, protection, shelter, evacuation, relocation, caution, damage, evacuation, injury, help, sympathy	damage, loss, economic, cost, impact, financial, property, disaster, recovery, reconstruction, insurance, business, investment, job_loss, economic_growth, market_disruption, supply_chain, infrastructure, resilience, recovery_funds, economic_development, employment, gdp_impact, financial_aid, bailout, debt, bankruptcy, taxation	dry, snow, high temperature, wind, thunderstorms, rain, winter, cold, summer, hot, lost moisture, pressure, water vapor, sea level pressure, precipitation

mate news articles, with topic keywords provided to climate data experts, to facilitate the development of the taxonomy. The taxonomy (Table B.3) organizes information into topics and subtopics, making it easy to explore extreme weather data from broad categories down to specific details.

Table B.3: A taxonomical organization of the analysis of extreme weather. Definitions of Extreme Weather Categories.

Topic	Definition
Vulnerabilities (V)	Alert communities based on their specific vulnerability factors
• Environmental Vulnerability (V)	Areas prone to greater damage due to fragile ecosystems, with location mention
• Infrastructure Vulnerability (V)	Structural deficiencies that increase the risk of damage during extreme weather events, with location mention
• Economic Vulnerability (V)	The susceptibility of an economy to financial losses due to disasters
Impact (I)	Immediate and long-term effects of extreme weather events
• Deaths (I)	Fatalities caused by extreme weather events
• Infrastructure Damage (I)	Physical harm to buildings, roads, and other critical structures
• Economic Damage (I)	Financial losses incurred
• Homeless (I)	People displaced due to the destruction
Emergency Response (E)	Immediate actions taken by authorities, organizations, and communities
• Evacuation (E)	Organized movement of people to safety from threatened areas
• Community Support (E)	Assistance provided to affected populations by both local organizations and broader networks
• Emergency Services (E)	Immediate rescue and medical aid delivered during a crisis
• Communication Strategies (E)	Plans for conveying crucial information to the public and responders during emergencies

Appendix B.1. Prompt Details

Figure B.1 shows one-shot prompt for vulnerability/impact/emergency statement assessment. Similarly, Figure B.2 and Figure B.3 shows one-shot prompt for emotion analysis and topic/subtopic/keyword labeling.

Figure B.1: One-shot prompt for vulnerability/impact/emergency statement assessment

```
[Task Description:]
Given the sentence: "<sentence>", determine which categories it belongs to based on the definitions below. Assign a probability score between 0 and 1 to each category (strictly numeric values only), reflecting confidence that the sentence fits each label. Ensure the final probability for all the categories sum is exactly 1.
[Definitions of Categories]
Vulnerability: Describes conditions that make people or places prone to harm, including:
- Forecasts or warnings for specific locations about hazardous conditions (e.g., storm alerts, flood watches).
- Excludes: General weather forecasts without warnings, past rainfall amounts, or climate trends without immediate hazard warnings.
- Special case: Mentions of inches of rain count only if linked to a forecast predicting danger.
Impact: Describes strictly measurable consequences of extreme weather, such as:
- Number of casualties, injuries, financial loss, infrastructure damage, or economic impact.
- Excludes: Mentions of states of emergency, inches of rain, road closures, tree falling, or event cancellations without quantifiable impact or flight delays.
- Special case: Casualties always count as Impact, even if estimates (e.g., 'dozens injured').
Emergency: Describes urgent actions requiring immediate response, including:
- Evacuations, rescues, emergency shelters, or disaster response efforts.
- Excludes: State of emergency declarations without mention of direct emergency actions.
Others:
- Sentences about extreme weather without clear Vulnerability, Impact, or Emergency markers.
- Sentences mentioning multiple events without enough detail to fit a single category.
- States of emergency, road closures, or school closures without measurable damage or emergency actions.
[Example:]
Input: Krakow is struggling after heavy rainfall, with city officials offering sandbags to protect homes.
<think>
Explanation:
1. Vulnerability: Krakow is struggling after heavy rainfall, indicating the city's vulnerability to flooding. The mention of heavy rainfall implies a risk, suggesting that Krakow is susceptible to potential harm due to the weather conditions. This fits the Vulnerability category as it shows the city's increased risk of flooding.
2. Impact: While the sentence mentions city officials offering sandbags to protect homes, there is no specific mention of measurable impacts such as casualties, injuries, financial loss, or infrastructure damage. Therefore, this category is less relevant compared to Vulnerability and Emergency.
3. Emergency: The fact that city officials are offering sandbags to protect homes implies an urgent response to a disaster, categorizing the action as an emergency. The immediate need for protective measures points to an emergency situation.
4. Others: This sentence describes a specific event related to heavy rainfall, but it is clearly categorized into Vulnerability, Impact, and Emergency. Therefore, it doesn't fit into the "Others" category as it clearly involves measurable consequences and urgent actions.
</think>
<output>
Final Output:
- Vulnerability: 0.40
- Impact: 0.10
- Emergency: 0.50
- Others: 0.00
</output>
```

Appendix B.2. Annotation Guidelines

This section outlines the detailed annotation process for dataset and gold testset preparation.

Appendix B.2.1. Dataset Preparation

To construct a high-quality alignment dataset tailored to extreme weather event analysis, we begin with the **ExtremeWeatherNews** corpus, which contains sentence-level reports from various news sources. The goal of this stage

Figure B.2: One-shot prompt for emotion analysis.

```
[Task Description]
Given the sentence: "<sentence>", determine which emotions it conveys. Assign probability scores ensuring they sum to exactly 1.
[Emotion Definitions]
Sadness: Expresses sorrow, grief, disappointment, or loss, often involving suffering or destruction.
Anger: Indicates frustration, outrage, or dissatisfaction, including expressions of blame or criticism.
Fear: Suggests concern, anxiety, or perceived threat, often linked to warnings or potential dangers.
Joy: Reflects happiness, relief, celebration, or positive outcomes.
Optimism: Shows hope, encouragement, or confidence in a positive future.
Trust: Indicates reliability, assurance, or faith in a person, system, or institution.
Neutral: Lacks strong emotional cues, purely factual, informational or descriptive.

[Example:]
Input: "The wildfire destroyed thousands of homes, leaving families devastated."
<think>
Explanation:
- The sentence describes destruction and suffering, strongly aligning with Sadness (0.8).
- There is minor frustration in the situation, leading to Anger (0.05).
- The threat aspect contributes to Fear (0.05).
- It lacks positive emotion, optimism, or trust.
- Since the sentence is descriptive but emotionally charged, Neutral (0.1) remains minimal.
</think>
<output>
Final Output:
- Sadness: 0.8
- Anger: 0.05
- Fear: 0.05
- Joy: 0
- Optimism: 0
- Trust: 0
- Neutral: 0.1
</output>
```

is to ensure that only clear, relevant, and logically consistent samples are used for training and evaluation.

Relevance and Clarity Assessment. Annotators are first asked to rate each sentence on two dimensions: **clarity** and **relevance** to the context of extreme weather events.

- **Clarity** captures whether the sentence is grammatically sound, unambiguous, and interpretable without requiring additional context.
- **Relevance** evaluates whether the sentence pertains to vulnerabilities, emergencies, or impacts caused by extreme weather, or if it diverges from this core scope.

Sentences rated as *low relevance* or *unclear* are discarded from the dataset to improve the signal-to-noise ratio and reduce noise during training.

Explanation Verification. Following sentence-level filtering, we construct the ExtremeAlign dataset, a alignment data constructed by adding task-specific

Figure B.3: One-shot prompt for topic/subtopic/keyword labeling.

```

[Task Description:]
Given the sentence: "<sentence>", determine the most relevant topic, subtopic, and keywords based on the definitions below. Assign a probability score between 0 and 1 to each possible category under Topic and Subtopic (strictly numeric values only), ensuring the final probability sum is exactly 1
[Topic and Subtopic Definitions]
Vulnerabilities: Identifies risk factors that make communities more susceptible to damage.
- Environmental Vulnerability: Fragile ecosystems that increase disaster impact.
- Infrastructure Vulnerability: Structural weaknesses leading to heightened risk.
- Economic Vulnerability: Financial instability due to disasters.
Impact: Direct effects of extreme weather.
- Deaths: Fatalities resulting from disasters.
- Infrastructure Damage: Physical destruction of buildings and roads.
- Economic Damage: Financial loss due to extreme weather.
- Homeless: Displacement due to destruction.
Emergency Response: Immediate actions taken to mitigate disasters.
- Evacuation: Organized movement of people to safety.
- Community Support: Aid provided by local and national organizations.
- Emergency Services: Rescue and medical response efforts.
- Communication Strategies: Plans for disseminating critical information.
[Example:]
Input: "Hurricane Maria devastated Puerto Rico, leaving thousands homeless and without power."
<think>
Explanation:
- The sentence describes significant destruction, categorizing it under Impact (0.8).
- "Leaving thousands homeless" aligns with Homeless (0.5), and "without power" suggests Infrastructure Damage (0.3).
- Emergency efforts likely followed, so Emergency Response (0.2) is relevant, particularly Emergency Services (0.2).
</think>
<output>
Final Output:
- Topic: Impact (0.8), Emergency Response (0.2)
- Sub-Topic: Homeless (0.5), Infrastructure Damage (0.3), Emergency Services (0.2)
- Keywords: devastation, homeless, power outage, Hurricane Maria
</output>
- Keywords should not include specific locations.

```

prompts using the sentences in ExtremeWeatherNews dataset aimed at training models for reasoning-aware alignment. Each sample in ExtremeAlign includes a sentence, a task-specific output, and an explanation chain justifying the output assignment.

To ensure consistency and interpretability, annotators are instructed to:

1. Read both the sentence and its associated explanation.
2. Verify whether the explanation logically and accurately supports the given task-specific output.
3. Remove or revise samples where the explanation is incomplete, inconsistent, overly generic, or incorrect based on domain definitions. In particular, explanations must align strictly with the formal definitions of categories.

Appendix B.2.2. Gold TestSet Annotation

To ensure high-quality and consistent annotations for the gold-standard evaluation set, we developed detailed guidelines covering each task in the

ExtremeAlign framework: (1) Vulnerability/Impact/Emergency Assessment, (2) Topic/Subtopic Labeling and Keyword Extraction, and (3) Emotion Analysis. These guidelines were designed to balance both explicit task definitions and the flexibility needed to incorporate human reasoning and domain knowledge.

1. General Instructions: Annotators were provided with a set of instructions for each task, including definitions, category descriptions, and examples of edge cases. To minimize ambiguity, we included both inclusion and exclusion criteria for each label, as well as clarifications for frequently confused categories. An initial training round was conducted, followed by calibration sessions to align interpretations across annotators. We also ask them to utilize the general domain knowledge and commonsense reasoning relevant to extreme weather events while annotating.
2. Vulnerability/Impact/ Emergency Assessment: Annotators were instructed to assess whether a sentence describes a situation involving vulnerability (pre-existing conditions that increase risk), impact (measurable damage or consequences of an event), emergency (urgent calls for help or crisis situations). Overlapping categories were allowed, and annotators were asked to assign probability scores across the categories such that they sum to 1. Reasoning for each choice was recorded in free-text format to capture implicit justifications.
3. Topic/Subtopic Labeling and Keyword Extraction: Each sentence was evaluated to determine its thematic focus based on a predefined taxonomy (Table B.3). Annotators could select multiple relevant topics and subtopics, and assign confidence scores to reflect relative importance. Annotators were encouraged to extract salient keywords that captured the essence of the sentence.
4. Emotion Analysis: Annotators labeled each sentence with one or more emotions from a defined set (Sadness, Anger, Fear, Joy, Optimism, Trust, Neutral). Probability distributions over the emotion labels were allowed to reflect emotional ambiguity.
5. Quality Control: Annotations were reviewed by a second annotator, and disagreements were adjudicated by a third expert.

Annotator Details and annotator agreement. Our annotators and verifiers comprise postdoctoral researchers and senior Ph.D. students with domain expertise in both natural language processing (NLP) and extreme weather analysis. Their interdisciplinary background ensures a nuanced understanding of the linguistic and scientific aspects of the task. Each annotation is evaluated against a well-defined rubric to maintain consistency and reliability across the dataset. Verifiers conduct an additional round of quality checks, resolving ambiguities and refining annotations as needed to meet gold-standard quality. Inter-annotator agreement was measured using Fleiss Kappa score. We obtained an average score of 0.88 for various task showing high agreement between the annotators. Only samples with high agreement were retained in the final gold set.

Appendix B.3. ClimaEmpact Online Dashboard

The ClimaEmpact framework is accessible online, providing rapid and near real-time analysis of extreme weather events. The displayed snapshot showcases various locations impacted by Typhoon Yagi, along with associated impacts, vulnerabilities, topics, subtopics, keywords, and emotion analysis.

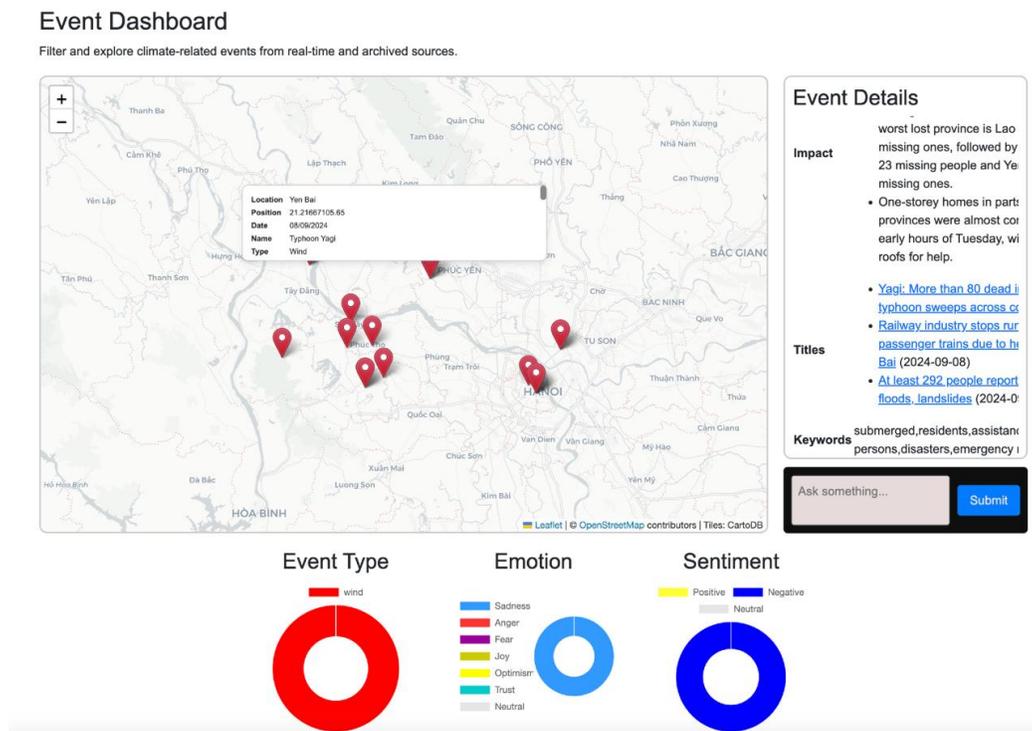


Figure B.4: ClimaEmpact Online Dashboard for Extreme Weather Analysis