AI Agency in Information Retrieval: Advancing Data Access Through Natural Language Interfaces

A Technical Report submitted to the Department of Computer Science

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Christopher Dale Barfield Jr.

Spring, 2024

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Rich Nguyen, Department of Computer Science

AI Agency in Information Retrieval: Advancing Data Access Through Natural Language Interfaces

Christopher Dale Barfield Jr.[a]
[a]University of Virginia School of Engineering, Department of Computer Science, Charlottesville, Virginia

Abstract. The integration of Natural Language Processing (NLP) interfaces into flood prediction systems aims to substantially enhance user experience by facilitating easy access to diverse data sources. Our project, Floodwatch, employs a multilingual agent, leveraging the LangChain framework atop a base NLP system and a large language model (LLM), such as GPT-4, to provide dynamic and adaptable interaction across various languages. This setup allows users to access real-time environmental data—from localized rain gauges to extensive weather APIs—and interact with machine learning predictive models that forecast flood probabilities. Initial findings from our implementation in Charlottesville indicate that users find the natural language interface significantly improves accessibility and utility of data. The system adeptly handles both specific queries and broad information requests, enhancing the relevance and precision of information retrieval. Importantly, the flexibility to incorporate additional data sources and execute machine learning inference on-demand promises extended capabilities and improved prediction intervals. However, the reliance on LLMs necessitates continuous monitoring to mitigate the risks of misinformation, particularly critical in disaster-related scenarios. These developments underscore the potential of NLP in transforming information retrieval practices in environmental monitoring systems, aligning technological advances with user-centric needs.

Keywords: artificial intelligence, natural language processing, information retrieval, human-computer interaction, conversational agents.

*Christopher Dale Barfield Jr., cdb8da@virginia.edu

1 Introduction

In disaster management, the timely and accurate delivery of data can be the difference between safety and catastrophe. Within the application floodwatch.io, a data and flood prediction platform, we are creating a new feature to enhance user interaction and data accessibility. Traditional methods within Floodwatch rely heavily on visual map interfaces and predetermined data updates, which are effective but have room for improvement in responsiveness and flexibility, especially across diverse linguistic landscapes.

The core of this innovation lies in integrating advanced natural language processing (NLP) into the Floodwatch platform. This integration addresses the critical need for improved accessibility and immediacy in data interaction. Our traditional system, while robust and useful, generally

create flood predictions and fetch related data in batch, with forecasts available up to three days in advance. This approach is complemented by tileset generation for regional visualization, but lacks the on-demand, user-driven interaction that enabled by modern forms of technology.

A significant challenge in expanding Floodwatch's utility across different linguistic demographics is the process of manual data translation, typically handled through the i18n framework, which is a widely used Javascript framework for localization of web applications. As Floodwatch grows to encompass more complex and varied data sources—including weather APIs, custom predictive models, and sensor technology like lidar and rain gauges, the need for a more dynamic and responsive communication interface becomes evident.

The introduction of an AI-driven NLP agent represents a pivotal development in making data from Floodwatch more accessible and actionable. This AI agent allows users to interact with the platform in their native language, requesting and receiving information seamlessly. The focus of this research and the subsequent development is not merely on enhancing flood prediction capabilities but on revolutionizing how users access, interact with, and utilize the vast array of data provided and used by the Floodwatch platform. This breakthrough is particularly beneficial for our global collaborators and users in Vietnam, France, Germany, and beyond as all new data integration is by default multi-lingual.

2 Background and Related Work

Floodwatch is at the forefront of integrating advanced interaction technologies into traditional flood prediction systems. The platform currently leverages Mapbox Tileset to visualize flood data on web-based maps, providing a robust but static user experience. The incorporation of a natural language processing (NLP) agent represents a strategic enhancement, aiming to transform the user

interaction from basic map navigation to dynamic, multilingual dialogue.

2.1 Existing Technologies

Traditional flood prediction systems often depend heavily on visual data representations such as those provided by Mapbox. These systems are effective for a broad overview but tend to be less flexible in responding to user-specific queries or providing real-time data updates. The shift towards incorporating NLP within Floodwatch is designed to complement these visual tools by offering a more interactive and responsive user experience, where data can be queried and presented through natural language.

2.2 Literature Review

While this project is pioneering in its specific application, it aligns with broader research trends emphasizing the efficacy of multilingual large language models (LLMs). The study "Multilingual LLMs are Better Cross-lingual In-context Learners with Alignment" underpins the theoretical framework for Floodwatch's NLP feature. This research highlights the potential of LLMs to enhance cross-lingual understanding and interaction, which is pivotal for Floodwatch as it expands to serve a global user base with diverse linguistic needs.

2.3 Comparative Analysis

Unlike traditional flood prediction interfaces that primarily offer graphical data interaction, Floodwatch's NLP agent introduces a novel, user-centric approach. This enhancement does not replace the existing functionalities but instead offers an alternative mode of interaction that is inherently flexible and accessible. By integrating multilingual capabilities directly into the NLP agent, Flood-

watch addresses the linguistic barriers that often complicate the use of such technologies in diverse regions.

2.4 Innovation in Data Sources

The integration of various data sources—ranging from localized sensor data to broader weather APIs—into Floodwatch is crucial for providing contextually rich responses to user queries. These sources feed into both the predictive models and the NLP interface, allowing the system to offer nuanced insights into weather conditions and flood risks. This approach not only enhances the predictive accuracy but also elevates the overall user experience by providing detailed, context-specific information that goes beyond generic flood predictions.

This background lays a solid foundation for understanding the significance of the enhancements being introduced to Floodwatch. By situating the project within current technological and research trends, it becomes clear how these innovations are set to advance the state of flood prediction and crisis management technology.

3 System Overview

The Floodwatch application operates as a PWA, delivering seamless user experiences via a tileset UI that visualizes flood predictions. These predictions originate from a machine learning (ML) predictive algorithm, specifically a Cartesian flood predictor, which is central to the platform's flood risk assessment capabilities. The recent integration of a natural language processing (NLP) feature aims to revolutionize user interactions by enabling an intuitive, conversational interface for accessing and querying flood-related data.

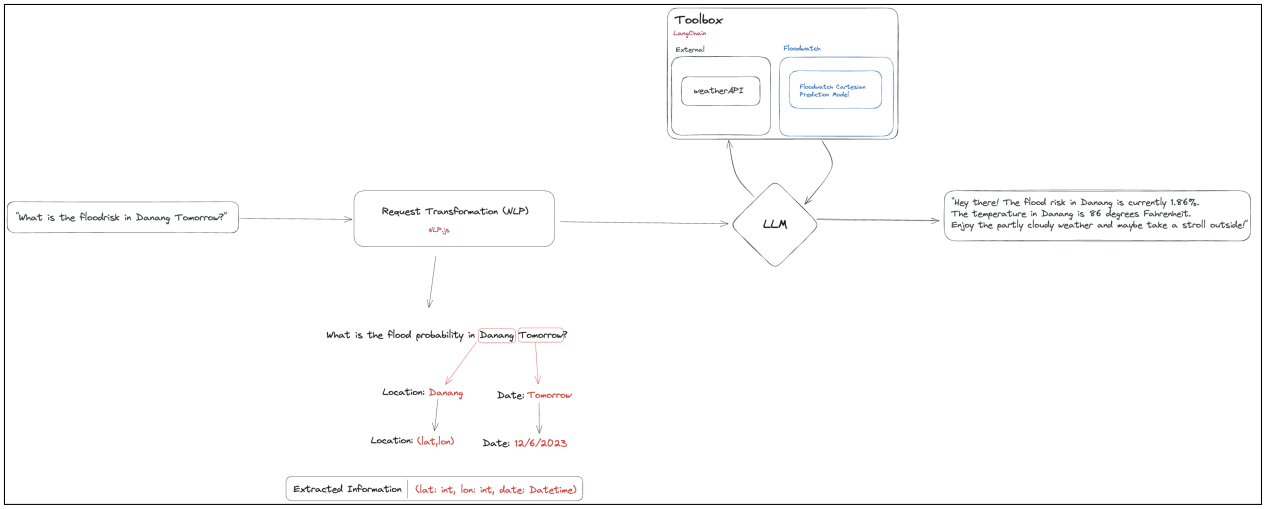

Fig 1 System Architecture of Floodwatch's NLP Agent. This diagram illustrates the overall structure of the Floodwatch application, highlighting the integration of the natural language interface, real-time weather data retrieval from APIs, and the interactive map display.

3.1 Natural Language Processing Feature

The NLP component within Floodwatch starts with the use of NLP.js for initial user query processing, ensuring that the system can accurately parse and understand a wide range of user inputs. After this initial processing, a more advanced Large Language Model (LLM) takes the lead. This LLM, capable of sophisticated data handling, retrieves necessary information from what Floodwatch terms as a 'toolbox'—a collection of integrated data sources including weather APIs and the ML predictive model. This allows the LLM to not only fetch relevant data but also run real-time inference based on user requests, culminating in a response crafted in natural language that is both informative and easy to comprehend.

The flexibility of this NLP feature is a significant enhancement over traditional interfaces. It allows users to ask complex questions about flood risks and weather conditions and receive tailored, specific information without needing to navigate through traditional data presentation formats. This feature supports multilingual interaction, enabling Floodwatch to serve a diverse global audience effectively and makes the platform highly accessible and user-centric.

3.2 Integration of Data Sources

In Floodwatch, data integration is meticulously managed to ensure the availability of timely and relevant data for accurate predictions and responses. This process is facilitated through direct connections to weather APIs that supply comprehensive weather data. This data is seamlessly piped into the flood prediction algorithm, enabling the system to conduct real-time inference based on up-to-the-minute environmental conditions.

3.3 Real-Time Processing

Transitioning from batch processing to real-time data handling, the NLP feature in Floodwatch introduces an on-demand processing capability. This innovation ensures that user queries trigger immediate system responses, with direct inference executed to provide current and accurate flood predictions and related information. This real-time processing capability significantly enhances the responsiveness of Floodwatch, crucial during rapidly evolving flood events.

Overall, the architecture of Floodwatch, enriched with its advanced NLP capabilities and robust data integration, sets a new benchmark for interactive and user-centric flood prediction systems. This pioneering approach not only elevates the user experience but also enhances the accessibility and precision of vital flood-related information, establishing Floodwatch as an indispensable tool in disaster management and response planning.

4 Evaluation

The evaluation of the newly implemented natural language processing (NLP) interface in Floodwatch focused primarily on user experience and the qualitative impact of the feature. This section

details the methodology used for evaluation, insights from initial user feedback, and observations from hands-on use.

4.1 Evaluation Methodology

Given the conversational and user-focused nature of the NLP feature, the evaluation was primarily qualitative. A small group of five users was selected to interact with the system, providing a mix of subjective feedback and observable data on their interaction patterns. This approach allowed for an in-depth understanding of how the feature was received by end-users and how it enhanced their experience with Floodwatch.

4.2 Comparison Metrics

While no formal quantitative metrics were established prior to this informal trial, the feedback received highlighted a significant improvement in user satisfaction due to the conversational nature of the interface. Users appreciated not only the ability to receive complex flood-related data but also the conversational elements that made the interactions feel more natural and engaging. For instance, suggestions to ”go outside and enjoy the weather” or reminders to ”bring an umbrella” if rain was expected were particularly well-received. These elements demonstrate the NLP interface’s capability to handle a larger set of information dynamically, making the system more responsive and user-friendly compared to traditional, static methods.

4.3 User Feedback

Feedback was collected through direct observation and informal interviews as users interacted with the NLP feature. This process allowed for real-time assessment of the interface’s performance and immediate user reactions. Observations focused on how users navigated the interface, the types

of queries they made, and their responses to the system's prompts and information delivery. The overall feedback was positive, with users noting the enhanced accessibility and usability of flood-related data through the NLP interface.

Without any formal case studies, the user sessions served as mini-case studies in themselves, providing valuable insights into how different types of users—varying in age, comfort with technology, and familiarity with the subject matter—interact with and benefit from the NLP feature.

4.4 Overall

The initial evaluation of the NLP interface in Floodwatch suggests that integrating conversational AI significantly enhances user engagement and satisfaction. It supports a more intuitive and accessible way for users to interact with complex environmental data. Although further studies with more rigorous quantitative metrics could provide deeper insights, the preliminary feedback is promising. This indicates a strong potential for broader implementation and further development of AI-driven interfaces in environmental data applications.

5 Results

The evaluation of Floodwatch's natural language processing interface provided insightful results, reflecting user interactions and system performance. Users expressed a genuine interest in the new feature, noting its innovative approach to information retrieval, although they acknowledged a learning curve compared to traditional map navigation. The ability to receive information in a conversational tone was particularly well-received, enhancing user engagement and making the data feel more accessible and personalized.

5.1 Performance of the Natural Language Interface

The NLP interface demonstrated a significant improvement in response times, with the AI capable of delivering answers in less than a second. This is a stark contrast to the existing batch processing and tileset generation methods used in Floodwatch, which are more resource-intensive and less immediate. The NLP feature's ability to perform direct inference on user-specified locations allows for more precise and tailored information, which could be especially beneficial for cross-lingual communication, aiding users in non-English speaking regions.

While the traditional map-based interface remains integral to the platform, the NLP interface offers a more dynamic and interactive alternative, potentially increasing the system's utility and user satisfaction. However, some users attempted to engage in follow-up questions, indicating a need for the system to handle more complex, conversational exchanges effectively.

5.2 Insights from User Experience

User feedback highlighted the value of real-time, conversational interactions provided by the AI. The system's setup allows the AI agent to access various tools (a 'toolbox') based on user requests, which currently includes a single weather API and the ML predictive model. As users showed interest in having more interactive and continuous dialogue capabilities, this points to an opportunity for further development in making the AI's conversations more context-aware and capable of handling follow-up interactions.

6 Discussion

6.1 Impact of AI Agency in Information Retrieval

The introduction of AI agency through the NLP interface in Floodwatch represents a significant step forward in making environmental data more accessible and interactive. This approach allows the AI to not only respond to user queries but also proactively offer advice and insights, enriching the user experience. However, as this feature evolves, it will be crucial to continuously monitor and refine the AI's responses to ensure accuracy and reliability, especially given the critical nature of the information being handled.

6.2 Implications for Future Developments

Looking ahead, the plan to enrich the AI's toolbox with more data sources and sophisticated ML models suggests a move towards a more robust, AI-driven approach to environmental monitoring. This strategy will likely set a precedent for similar applications in other fields, where real-time data processing and user-specific responses can significantly enhance decision-making processes.

6.3 Potential Expansions and Features

Future expansions will focus on enhancing the AI's toolbox and improving the user interface to support more complex interactions. Additional UI enhancements are planned to make the platform more intuitive and engaging. Expanding language coverage will also be a priority, ensuring that Floodwatch can serve a broader, more diverse global audience effectively.

These results and future plans underscore the transformative potential of integrating AI and NLP into traditional systems, promising not only to enhance user interaction but also to propel Floodwatch into a new era of technological advancement in disaster management.

7 Conclusion

This study set out to enhance the user experience of Floodwatch by integrating a natural language processing (NLP) feature powered by a large language model (LLM). The primary goal was to make data access more intuitive and efficient, allowing users to interact with the system through conversational AI. Our findings indicate that the LLM is capable of ingesting a large array of data and contextually filtering the necessary information to meet user requests. This capability is particularly significant as Floodwatch expands into diverse geographical regions, necessitating adaptive and accessible interfaces for global users.

The broader impact of AI in improving accessibility to information is profound. As users become more accustomed to interacting with natural language interfaces, the adoption of such technologies can enhance engagement and accessibility across various fields, not just in disaster management. This shift could redefine user expectations and experiences in interacting with complex data systems.

8 Acknowledgements

I would like to express my gratitude to my teammates working on the Floodwatch project for their collaboration and insights throughout this project. Special thanks are also extended to my principal investigator and professor, Rich Nguyen, who has provided invaluable guidance and support during this project. Their expertise and encouragement were instrumental in the successful integration of the NLP feature into Floodwatch.