MoE-Precip: Adaptive Mixture-of-Experts Fusion and an Interactive Web-Based Visual Analytics Interface for Multimodal Precipitation Forecasting

Quinn Krommenhoek
School of Computing and Design
California State University Monterey Bay,
Seaside, CA, 93955
qkrommenhoek@csumb.edu

Kofi Osei
Gustavus Adolphus College
St. Peter, MN, 56082
okosei@gustavus.edu

Sai Deepthi Yeddula
School of Engineering and Computer Science
Oakland University
Rochester, MI, 48309
syeddula@oakland.edu

Dongji Feng*
School of Computing and Design
California State University Monterey Bay
Seaside, CA, 93955
dfeng@csumb.edu

Chen Jiang†
Samuel Ginn College of Engineering
Auburn University
Auburn, AL 36849
czj0042@auburn.edu

Abstract

Accurate precipitation forecasting is essential for agriculture, disaster response, and climate adaptation, yet it remains challenging because the information needed for prediction is inherently multimodal. Radar, satellite imagery, and surface gauges provide complementary observations, but they differ substantially in spatial and temporal resolution and exhibit distinct domain characteristics. For example, atmospheric variables are similarly heterogeneous: u/v wind captures momentum, whereas temperature and humidity encode the thermodynamic state of the atmosphere. Nevertheless, many existing methods [1, 2] collapse these diverse inputs into a single shared representation, an approach that can be brittle and may generalize poorly across weather regimes and sensing conditions. At the same time, the growing complexity of multimodal forecasting models makes their behavior more difficult to interpret, thereby limiting users' ability to understand why a prediction succeeds or fails across regions, timesteps, and weather regimes.

To address this limitation, we present MoE-Precip, a visual analytics system that bridges complex multimodal precipitation modeling and practical decision-making. The system is built around an adaptive Mixture-of-Experts (MoE) architecture [3], in which expert subnetworks are organized by Earth-system components and sensing modalities, such as momentum and thermodynamics [4]. This structure enables specialized modeling of recurring spatio-temporal regimes, including convective bursts and rainband evolution. A dynamic router further learns context-dependent expert weights, allowing heterogeneous signals to be fused adaptively rather than compressed into a single shared representation.

*Corresponding author: dfeng@csumb.edu.
†Corresponding author: czj0042@auburn.edu.

To support interpretation and analysis, MoE-Precip provides a web-based interface through which users can scrub through time, select input modalities and features, and examine synchronized Observed–Predicted–Error views. These linked views allow users to rapidly identify forecast successes and failures, detect grid-level and regional biases, and diagnose breakdowns associated with specific weather regimes. The interface also summarizes error patterns across space, time, and regimes, helping users calibrate trust in the model and gain insight into its performance characteristics.

We tested MoE-Precip on a curated multimodal dataset[3] that captures real-world conditions during Hurricane Ian (2022) and its impacts in South Florida [5].

1 Introduction

Accurate precipitation forecasting is essential for informing agriculture and disaster management. Modern weather forecasting systems increasingly rely on integrating diverse, multi-source observational data, such as radar reflectivity, satellite imagery, and atmospheric measurements. The multi-source data provides invaluable insights into evolving weather patterns and opens new avenues for real-time monitoring and management. The trade-off in modeling with multi-source data is the data's highly heterogeneous nature. For example, data from multiple sources can be highly diverse, with varying spatial and temporal resolutions and physical features. This is the main reason that integrating multi-source data into commonly used deep learning architectures may be less successful than using single-source data.

Although recent advances in machine learning have improved weather prediction capabilities [1, 2], many models remain limited in their ability to efficiently process heterogeneous multimodal data inputs. In particular, variations in data format, temporal alignment, and regional relevance complicate joint learning across modalities. The standard architecture in the most commonly used modeling methodologies often treats all inputs uniformly, neglecting the fact that different data sources may contribute differently depending on the spatial or temporal context.

To address this limitation, we propose a visual analytics system that supports the exploration and analysis of climate-related variables while forecasting precipitation rates, providing a useful reference for severe weather monitoring and assessment. The system is centered on an adaptive Mixture-of-Experts (MoE) model for precipitation forecasting, in which each expert specializes in a specific data modality or spatiotemporal pattern. Then, a dynamic router routes input to the most relevant experts according to the current context. This modular architecture achieves a balance between specialization and scalability, improving both predictive accuracy and interpretability under complex climate conditions. The system also features a browser-based visualization Interface that allows users to interactively inspect climate variables used in the prediction model. Through transparent views of weather patterns across space and time, the platform facilitates a deeper understanding of model behavior and supports more effective analysis of severe weather events.

2 System Design

2.1 MoE-Climate Dataset

This section introduces the *MoE-Climate* dataset used in this project, and it supports a variety of climate-related modeling tasks. The raw data used to construct the MoE-Climate dataset were collected by the National Oceanic and Atmospheric Administration (NOAA) using multiple remote-sensing platforms and in-situ devices. These sources vary in both spatial and temporal resolution and collectively cover the entire continental United States. For this study, we extracted and processed a regional subset centered on South Florida to capture the atmospheric dynamics of a specific extreme weather event: Hurricane Ian.

As shown in Table 1, this dataset includes high-resolution gridded climate observations for the South Florida region during the landfall of Hurricane Ian in 2022 [5]. The temporal span extends from 00:00 on September 23 to 23:00 on October 1, with an hourly temporal resolution. The data is spatially organized into a 100×100 grid, where each grid cell corresponds to a 3 km $\times$ 3 km area, allowing fine-grained monitoring of localized weather patterns.

Each point in the grid includes a comprehensive set of atmospheric variables sampled at 50 vertical pressure levels. These variables include precipitation rate, total accumulated precipitation, cloud cover, and 18 additional climate-related features. As detailed in Table 1, the dataset also provides metadata, including grid id, latitude, longitude, global grid dimensions (1799×1059), and grid spacing. Each observation is timestamped, allowing for dynamic temporal analysis throughout the entire hurricane period.

[3] `https://www.noaa.gov`

Table 1: A snapshot of the climate data used in the study that contained both spatial and temporal inputs. The dataset included grid metadata and time-series measurements, including precipitation rate, total precipitation, cloud cover, and 16 additional meteorological features.

GridID	Longitude	Latitude	Grid Points	Grid Spacing	Vertical Level
1	122.71	21.13	1799 × 1059	3 km	50

Time Stamp	2022/09/23 00:00	2022/09/23 01:00	2022/09/23 02:00	...	2022/10/02 00:00
Precipitation rate (mm/hour)	0.0	0.72	0.94	...	0
Total Precipitation (mm)	0.01	1.88	4.3	...	31.61
...	...	...	...	...	...
Cloud cover	50%	70%	100%	...	20%

2.2 Knowledge-Guided Feature Grouping

To support the design of our modular MoE model, we partitioned the full set of climate variables into subsets according to their shared physical properties. This knowledge-guided grouping is determined by domain-specific classifications provided by the National Oceanic and Atmospheric Administration (NOAA) [4]. As shown in Table 2, the variables are organized into six physically meaningful categories: Momentum, Temperature, Moisture, Mass, Cloud, and Radiation.

2.3 System Architecture

As shown in Figure 1, the MoE-Precip system consists of two main components: a multimodal precipitation prediction model and an interactive Visualization and Analytics Interface. The overall design connects predictive modeling with human-centered visual analysis.

The prediction model integrates heterogeneous weather data from multiple sources. These multimodal inputs are first fed into a shared latent space, which aligns information from different data sources into a unified representation. This shared feature space supports joint learning across modalities while preserving complementary weather information.

Based on this representation, the model employs an MoE architecture with specialized experts for data collected from different sources. Each expert is designed to focus on a physically meaningful subset of variables. An adaptive gating network then dynamically weights the contributions of these experts according to the current weather context, and the fused expert outputs are used to produce the final precipitation prediction.

The Visualization and Analytics Interface complements the model by allowing users to explore both input variables and forecast results. The interface supports visual inspection of weather conditions such as cloud cover, moisture, and temperature, as well as comparisons of observed and predicted precipitation and the prediction error. This side-by-side design helps users identify spatial patterns, assess forecast quality. The technical details of the interface are presented in Section 3.

3 Visualization Interface System

3.1 System Overview

Climate data, with their high density and multi-dimensionality, are inherently difficult to interpret directly. Our goal is to enable seamless transitions across different data representations, allowing us to better interpret our model's predictions by visualizing how data points are distributed across space and time. To achieve this goal, a Web-Based Application was developed as a thin layer of abstraction to reduce cognitive load when working with a broad, heterogeneous dataset. When observed in real-time, it is important that the subtle nuances in the changing data can be assessed with precision. Data is read directly from the file and indexed by date and time. Each point is then mapped to its corresponding 3km x 3km point and time index. This approach has been optimized to minimize overhead and update quickly so that the broad geographical range can be easily observed over time. The data set provided consists of 10,000 grid points every hour over 9 days, which is hardly legible to the human eye in its traditional format.

3.2 Technical Framework

Our web application was made using HTML5, CSS3, and standard JavaScript up to date with ES2025 standards. For portability, the application is implemented as a static web application that requires only an HTTP file server for deployment. The rendering is powered by Leaflet.js [4] a well-established open-source library. The map layouts consist

[4] `https://leafletjs.com/`

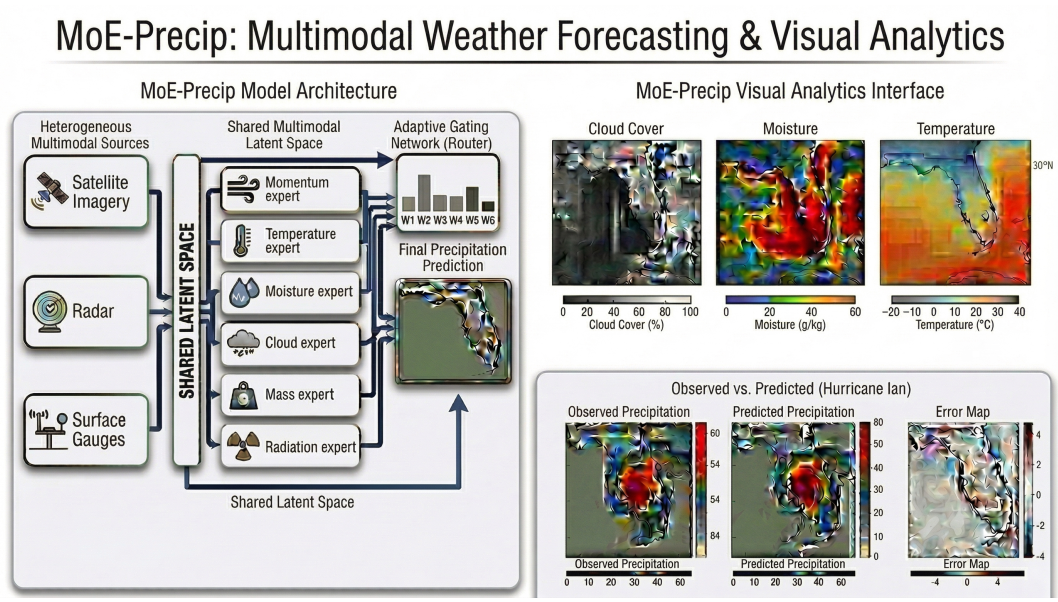

Figure 1: Overview of the MoE-Precip framework, including the multimodal forecasting architecture and the visual analytics interface, this figure is shown for illustrative purposes.

of an OpenStreetMap[5] for street view, ArcgIsOnline[6] for Satellite, and OpenTopoMap[7] for topography. JavaScript has become a staple in the industry, leading to most major web browsers supporting it. The SpiderMonkey(FireFox) and V8(Google) engines have trade offs, but have both made JavaScript run more efficiently in the browser with the use of WebAssembly. The emphasis of minimizing dependencies eliminates the problem of out of date frameworks. The universal use of JavaScript leads to frameworks for streamlined processes such as React. However, recent history has shown that applications dependent on this infrastructure can pose security threats[8]. The application runs on a server hosted by California State University, Monterey Bay.Our code can be found in GitHub

3.3 Visualization Design

The Six categories described in Table 2 vary in the units that represent them and in their respective numerical range. The method of representing this was simply to apply a rainbow scale [6, 7] which represents lower or weaker numbers as more blue and becomes more red as the value increases. Due to the varying input ranges, when looking at the visualization for anything other than temperature, the subtle variations are less discernible. After presenting and discussing this with HCI experts, it was found that specialized color mapping provides an accurate representation of the data.

For instance, in Figure 2d, we can observe patterns common to temperature, such as a cold front denoted by the large blue area. In contrast, for cloud coverage you can determine where heavy cloud formations form. Using a blue-scale for precipitation and related fields substantially improved interpretability, while fields with an easily identifiable range, such as cloud cover, are well represented as Figure 2a.

4 Implementation

When loaded, the visualization creates a color range to represent possible values. Each color is determined by its position within the range of values. This is intended to present the information in a more familiar manner, making comprehension more intuitive. Contemporary forecasting systems use the rainbow scale as it easily expresses extreme differences in the data, but for forecasts not representing temperature it is more intuitive to have a scale that represents rain rather than general meteorological intensity. To better understand how variables change over time, an option is available to automatically scrub through the time index.

[5] `https://openstreetmap.us/`
[6] `https://www.arcgis.com/index.html`
[7] `https://opentopomap.org/`
[8] `https://react.dev/blog/2025/12/03/critical-security-vulnerability-in-react-server-components`

Table 2: Feature group mapping. Specific features are referred to as their corresponding Feature ID in the released source code.

Feature Group	Features	Feature ID
Momentum	Wind speed	5
	U component of wind	8
	V component of wind	9
Temperature	2 metre temperature	6
	Dew point temperature	10
Moisture	Precipitation rate	1
	Plant canopy surface water	3
	Moisture availability	11
	Total precipitation	12
	Relative humidity	16
Mass	Pressure:CloudBase	7
	Pressure:CloudTop	15
Cloud	Cloud cover	2
	Low cloud cover	13
	Medium cloud cover	14
Radiation	SBT collected by GOES 11 C3	18
	SBT collected by GOES 11 C4	19
	SBT collected by GOES 12 C3	4
	SBT collected by GOES 12 C4	17

SBT stands for Surface Brightness Temperature, a satellite-derived radiation variable included under the Radiation group; GOES stands for Geostationary Operational Environmental Satellite; C stands for channel.

4.1 Prediction Analysis

The multi-dimensional nature of the data requires human inquiry and reasoning. The application offers a "Prediction View," shown in Figure 3, that renders predicted values alongside the actual values. The prediction view updates simultaneously with the actual view, showing the values side by side. The legend remains the same for both predicted and actual values to ensure consistent color differences. The predicted view is consistent with the bounds and coloring of the actual map. Hovering over a grid cell on any map simultaneously displays the corresponding values. In addition, we implemented an Error Map View to explicitly visualize the difference between predicted and actual values. Together, these views retain the main characteristics and constraints of the original map design, promoting an honest and effective comparison.

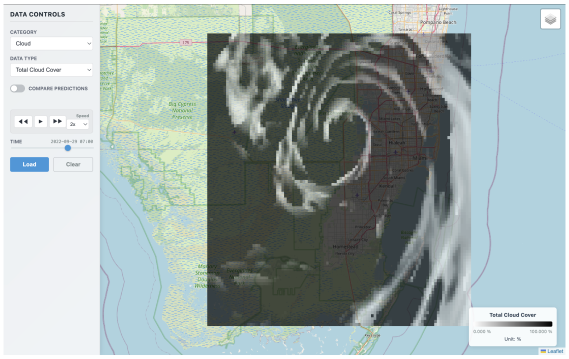

(a) Total Cloud Cover

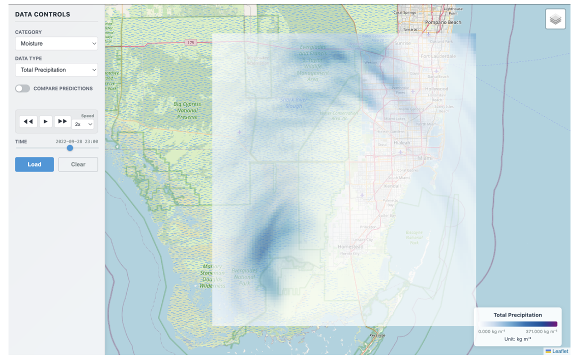

(b) Total Precipitation

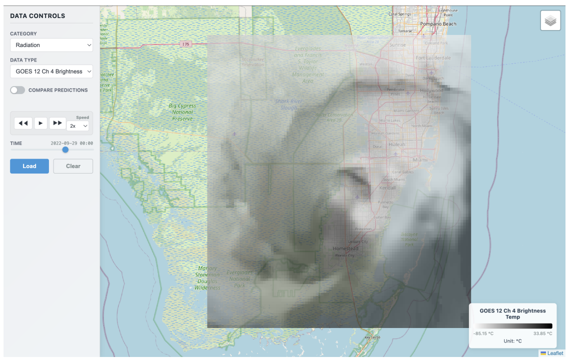

(c) GOES-12 Channel 4 Brightness Temperature

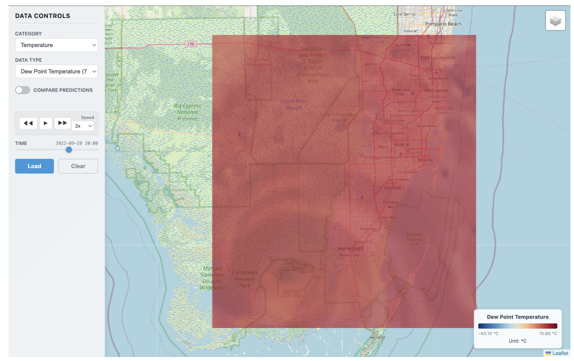

(d) Dew Point Temperature

Figure 2: Interactive Visualizations of the dataset over South Florida. (a) displays widespread cloud presence. (b) displays total precipitation, illustrating localized rainfall patterns with areas of higher accumulation. (c) depicts GOES-12 Channel 4 brightness temperature, where colder values indicate higher cloud tops associated with convective activity. (d) shows spatial distribution of dew point temperature, highlighting consistently high moisture levels across the region

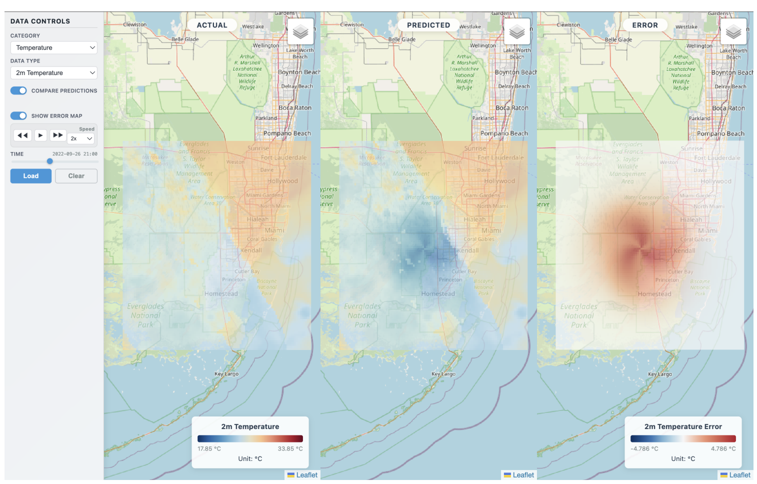

Figure 3: Interactive visualization of 2 m air temperature over South Florida at 21:00 on September 26, 2022, showing the actual temperature observations (left), model-predicted temperature (center), and the corresponding prediction error map (right). The interface supports side-by-side comparison of observed and predicted spatial temperature patterns.

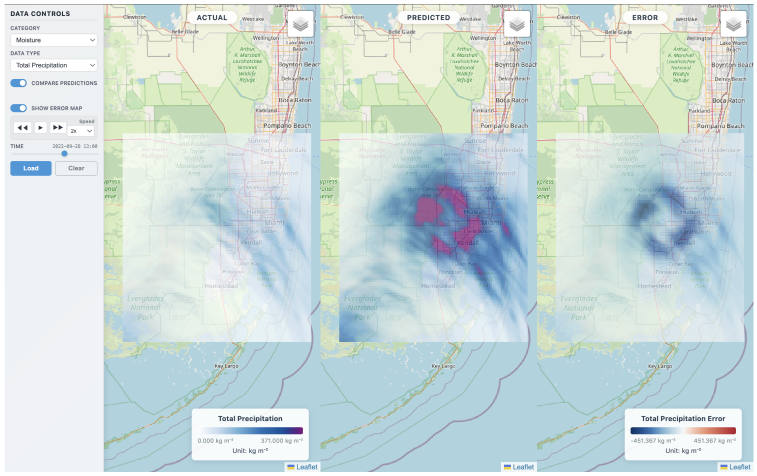

Figure 4: Interactive visualization of total precipitation over South Florida at 13:00 on September 28, 2022, showing the actual observations (left), model-predicted precipitation (center), and the corresponding prediction error map (right). The interface allows side-by-side comparison of spatial precipitation patterns and highlights areas of overestimation and underestimation.

5 Conclusion and Future Work

This application reduces cognitive load by allowing users to view complex meteorological data across space and time. Because geoscientific prediction involves large volumes of heterogeneous variables that evolve continuously, understanding model behavior requires more than examining isolated values from the dataset. Through the design of coordinated views, consistent color mapping, and side-by-side comparisons of observed, predicted, and error maps, the interface helps users to identify where and when forecasts succeed or fail.

This interactive application can also be extended to other spatiotemporal tasks [8, 9], including traffic prediction, crop yield prediction, and wildfire detection. In future work, we plan to broaden the scope of the interface to support a wider range of applications and more informed decision-making.

6 Acknowledgements

We thank California State University (Monterey Bay), School of Computing and Design, for their continuous support.

References

[1] Afan Galih Salman, Bayu Kanigoro, and Yaya Heryadi. Weather forecasting using deep learning techniques. In *2015 international conference on advanced computer science and information systems (ICACSIS)*, pages 281–285. Ieee, 2015.

[2] Yoo-Geun Ham, Jeong-Hwan Kim, and Jing-Jia Luo. Deep learning for multi-year enso forecasts. *Nature*, 573(7775):568–572, 2019.

[3] Saeed Masoudnia and Reza Ebrahimpour. Mixture of experts: a literature survey. *Artificial Intelligence Review*, 42:275–293, 2014.

[4] NOAA National Centers for Environmental Information (NCEI). Data access, 2024. Accessed June 12, 2025.

[5] Chen Jiang, Wenlu Wang, Naiqing Pan, and Wei-Shinn Ku. A multimodal geo dataset for high-resolution precipitation forecasting. In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems*, pages 1–4, 2023.

[6] Yang Liu and Jeffrey Heer. Somewhere over the rainbow: An empirical assessment of quantitative colormaps. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pages 1–12, 2018.

[7] BRITTNEY Bryant, Matthew Holiner, RACHAEL Kroot, KATHLEEN Sherman-Morris, WILLIAM B Smylie, LISA Stryjewski, MEAGHAN Thomas, and CHRISTOPHER I Williams. Usage of color scales on radar maps. *J. Oper. Meteor*, 2:169–179, 2014.

[8] Sai Deepthi Yeddula, Chen Jiang, Bo Hui, and Wei-Shinn Ku. Traffic accident hotspot prediction using temporal convolutional networks: A spatio-temporal approach. In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems*, pages 1–4, 2023.

[9] Chen Jiang, Sai Deepthi Yeddula, and Wei-Shinn Ku. A convolutional neural network model for accurate short-term leaf area index prediction. *Modeling Earth Systems and Environment*, 10(5):6397–6412, 2024.