



Research article

Exploring large language models for climate forecasting

Yang Wang* and Hassan A. Karimi

Geoinformatics Laboratory, School of Computing and Information, University of Pittsburgh, 135 N Bellfield Ave, Pittsburgh, PA 15213, USA

* Correspondence: Email: yaw70@pitt.edu; Tel: +14127086722.

Academic Editor: Pasi Fränti

Abstract: With the increasing impacts of climate change, there is a growing demand for accessible tools that can provide reliable future climate information to support planning, finance, and other decision-making applications. Large language models (LLMs), such as GPT-4o, present a promising approach to bridging the gap between complex climate data and the general public, offering a way for non-specialist users to obtain essential climate insights through natural language interaction. However, an essential challenge remains underexplored: Evaluating the ability of LLMs to provide accurate and reliable future climate predictions, which is crucial for applications that rely on anticipating climate trends. In this study, we investigated the capability of GPT-4o in predicting rainfall at short-term (15-day) and long-term (12-month) scales. We designed a series of experiments to assess GPT's performance under different conditions, including scenarios with and without expert data inputs. Our results indicated that GPT, when operating independently, tended to generate conservative forecasts, often reverting to historical averages in the absence of clear trend signals. This study highlights the potential and challenges of applying LLMs for future climate predictions, providing insights into their integration with climate-related applications and indicating directions for enhancing their predictive capabilities in the field.

Keywords: large language models; climate prediction; rainfall

1. Introduction

As the impacts of climate change intensify, obtaining accurate information about future climate trends has become increasingly important. Fields such as energy planning, urban development, and weather derivatives increasingly rely on precise climate forecasts to make critical decisions [1]. However, accessing, analyzing, and interpreting climate data typically requires interdisciplinary knowledge in areas such as climatology, geography, statistics, and computer science, making it

challenging for the general public to utilize this information effectively. Thus, there is a growing need for straightforward, accessible ways for the public to obtain relevant and important information or services based on climate predictions [2].

In recent years, large language models (LLMs) like ChatGPT-4 have provided a convenient means for the public to access specialized information [3]. With advanced natural language processing capabilities, LLMs can provide individuals with complex domain-specific knowledge. By processing simple language queries, LLMs enable users to obtain the necessary insights without requiring any training [4]. The remarkable progress in knowledge integration and common-sense reasoning exhibited by LLMs has led to their exploration in various data-intensive fields, including scientific computing and financial forecasting, making climate science a promising area for LLM applications [5,6].

In the climate domain, several researchers have attempted to apply LLMs to professional communication and information dissemination [7,8]. For instance, the ClimSight project uses future forecast data from the Climate Modeling Intercomparison Project (CMIP) to provide agricultural recommendations through local climate services [9]. ClimateGPT, on the other hand, trains LLMs on climate science literature to generate more specialized climate knowledge and descriptive content [10]. However, most researchers investigating climate LLMs focus on “descriptive output”; in particular, how to make LLMs generate content that reflects climate science terminology accurately [11]. LLMs have also been applied to simulating public opinions and diverse perspectives in the context of science communication about climate change [12,13]. While these models excel at interpretative tasks, the essential component of LLM responses for future climate scenarios lies in their ability to accurately predict future climate trends. Without this capacity to predict future climate factors reliably, LLM outputs may fall short of meeting the practical demands of climate adaptation and planning.

To understand what the general public, who is not trained in climate modeling and analysis, would be able to get from LLMs on climate predictions, we aim to explore and assess the performance of LLMs in climate prediction tasks, focusing on their ability to capture trends when generating future climate data. We selected ChatGPT-4o, a representative LLM, for analysis and designed a series of experiments to evaluate its performance in generating short-term (15-day scale) and long-term (12-month scale) rainfall forecasts. GPT-4o is an improved LLM launched by OpenAI, optimized based on the GPT-4 architecture (with “o” representing “optimized”). GPT-4o excels in logical reasoning, long-document processing, and answering interdisciplinary questions. By enhancing semantic understanding capabilities and improving context memory mechanisms, it more efficiently supports tasks such as data analysis and scholarly translation. Moreover, GPT-4o demonstrates a deeper grasp of domain-specific knowledge. The contribution of our study lies in advancing beyond climate LLM studies that focus primarily on descriptive output. We explore the predictive capabilities of LLMs, specifically their potential to generate accurate future climate trends with and without the assistance of domain-specific knowledge.

The structure of this article is as follows: In Section 2, we introduce the data used and the model employed as the rainfall expert model (EM). In Section 3, we describe the experiments conducted. In Section 4, we present the comparative results. The discussion and conclusions are provided in Section 5.

2. Data and rainfall expert model

In this study, we focused on rainfall prediction at two time scales: short-term (15-day scale) and long-term (12-month scale). We considered 15 cities across the contiguous United States: Washington DC, Tucson, AZ, Salt Lake City, UT, Reno, NV, Phoenix, AZ, Pensacola, FL, New York, NY, Mobile,

AL, Forks, WA, El Paso, TX, Dallas, TX, Chicago, IL, Birmingham, AL, Baton Rouge, LA, and Atlanta, GA. These cities represent varying levels of rainfall, covering high, medium, and low rainfall areas. For each city, we collected daily data on maximum temperature, minimum temperature, and rainfall, which were used as inputs for future rainfall prediction. Figure 1 shows the locations of the cities and their annual average rainfall.

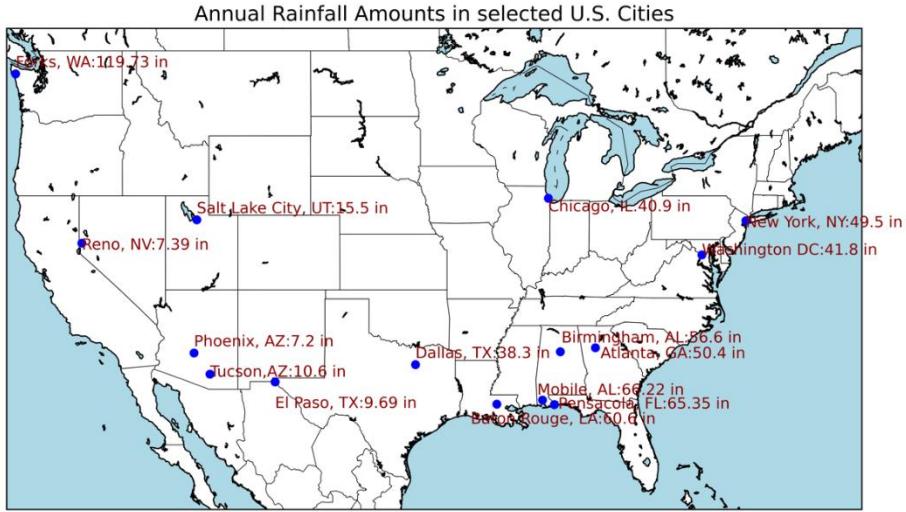


Figure 1. The locations of selected cities in the United States and their corresponding annual rainfall amounts.

For rainfall prediction, historical daily maximum temperature, minimum temperature, and rainfall were used as inputs. For temperature prediction, only minimum and maximum temperature were used as inputs. After preprocessing, the data from 1900 to 2022 were used, with 80% for training and 20% for validation.

We employed a two-layer LSTM model as our rainfall expert model (EM) for rainfall prediction. LSTM is a recurrent neural network (RNN) specifically designed to capture long-term dependencies in sequential data by mitigating the vanishing gradient problem through its unique gating mechanism [14,15]. The model used an input window of 60-time steps, corresponding to 60 days for short-term prediction and 60 months for long-term prediction. The output window was set to 15-time steps for short-term prediction and 12-time steps for long-term prediction, respectively. The LSTM hidden size was set to 128, and a batch size of 64 was used. The model was trained over 500 epochs, with the Adam optimizer. The specific input features for rainfall prediction included historical minimum temperature, maximum temperature, and rainfall over the input window, predicting future rainfall. For temperature prediction, only the corresponding minimum and maximum temperatures were used as inputs.

To evaluate the model's predictive accuracy, we used data starting from September 30, 2023, as the baseline for comparison. For short-term predictions, we used the 60 days prior to this date to forecast daily rainfall from October 1 to October 15, 2023. For long-term predictions, we used the 60 months prior to October 2023 to forecast monthly rainfall from October 2023 to September 2024. This setup enabled us to evaluate the EM model's performance on both short-term and long-term rainfall forecasting tasks.

3. Experimental setup

In this study, to systematically evaluate the performance of LLMs in rainfall forecasting tasks and

to analyze their ability to generate future climate data trends and patterns under input scenarios, we designed several experiments. These experiments were intended to assess the LLM's ability to generate rainfall predictions with and without the support of specialized knowledge. The different experimental conditions also ranged from direct to indirect provision of expert information, enabling a comprehensive analysis of GPT's potential in integrating climate data and generating forecasts. The selection of the most suitable prompt in this experiment depended mainly on intuition and experimentation. The experimental setups were as follows:

3.1. Experiment 1: GPT-only prediction

In the first experiment, GPT-4o was tasked with independently generating rainfall predictions without any additional expert information or data input. This condition mainly assesses GPT's ability to produce short-term and long-term rainfall predictions based solely on its pre-trained knowledge. By analyzing its predictions, we aim to uncover GPT's inclination in judging future climate trends without external guidance. This experiment helps us understand the GPT's interpretation of climate events based on its common-sense knowledge. We used the following prompt sample:

Prompt Sample 1:

You are a climate data prediction system focused primarily on forecasting rainfall for selected cities. Your timestamp is September 30, 2023, meaning you only consider information available prior to this date. Please make a final forecast based on your knowledge, including historical trends, regional variations, and potential future scenarios. For the time being, please ignore narrative responses; I am only interested in numerical results. Please predict for {city} during {October 1, 2023, to October 15, 2023}.

Please use the supplied data to predict the rainfall for the above period.

3.2. Experiment 2: GPT-EM prediction

In the second experiment, we provided GPT-4o with rainfall predictions generated by a rainfall expert model (EM) and asked it to make further predictions based on this expert data. This condition was designed to examine whether GPT, after receiving direct forecast support from a specialized model, can effectively utilize this data to capture trends or adjust its predictions. By comparing GPT's predictions with and without the expert data support, we could analyze its ability to integrate external information. We used the following prompt sample:

Prompt Sample 2:

You are a climate data prediction system focused primarily on forecasting rainfall for selected cities. Your timestamp is September 30, 2023, meaning you only consider information available prior to this date. I will provide you with a potential {daily} prediction for the period {October 1, 2023, to October 15, 2023} based on a deep learning model for the {city}. Please consider the results of the model and combine them with your knowledge to make a final forecast. For the time being, please ignore narrative responses; I am only interested in numerical results.

—potential forecast—

Period: {October 1, 2023 to October 15, 2023}

Rainfall: {v1, v2, v3, v4, v5, v6, v7, v8, v9, v10, v11, v12, v13, v14, v15}

Please use the supplied data to predict the rainfall for the above period.

3.3. Experiment 3: GPT-Regional climate prediction

In the third experiment, we indirectly provided GPT with predictions of regional climate factors related to rainfall (minimum and maximum temperatures) and tasked it with generating rainfall predictions based on these climate factors. In this setup, GPT did not receive direct rainfall predictions but instead relied on relevant climate variables as hints to infer trends. We aimed to evaluate whether GPT could use indirect climate information to generate reasonable rainfall predictions and assessed its ability to understand the relationships between climate variables in the absence of explicit rainfall data. We used the following prompt sample:

Prompt Sample 3:

You are a climate data prediction system focused primarily on forecasting rainfall for selected cities. Your timestamp is September 30, 2023, meaning you only consider information available prior to this date. I will provide you with a potential {daily} prediction for the period {October 1, 2023, to October 15, 2023} based on a deep learning model for the {city}. These predictions include {daily} maximum and minimum temperatures. Please consider the relationship between these climate data and potential rainfall. Integrate this information with your knowledge to make a final prediction. For the time being, please ignore narrative responses; I am only interested in numerical results.

—potential forecast—

Period: {October 1, 2023, to October 15, 2023}

Tmin: {v1, v2, v3, v4, v5, v6, v7, v8, v9, v10, v11, v12, v13, v14, v15}

Tmax: {v1, v2, v3, v4, v5, v6, v7, v8, v9, v10, v11, v12, v13, v14, v15}

Please use the supplied data to predict the rainfall for the above period.

3.4. Experiment 4: GPT-Teleconnection prediction

In the fourth experiment, to further explore GPT's understanding of large-scale climate indices, we provided it with global teleconnection indices, such as Nino3.4 [16], the Pacific Decadal Oscillation (PDO) [17], and the North Atlantic Oscillation (NAO) [18] and instructed it to generate a 12-month-scale rainfall forecast based on this predictive information [19]. These teleconnection indices are crucial variables in climate forecasting, revealing the long-term influence of large-scale climate patterns on variables such as rainfall [20,21]. It should be noted that in this experiment, we used the actual values of these indices from October 2023 to September 2024 as inputs to serve as a comparison. We used the following prompt sample:

Prompt Sample 4:

You are a climate data prediction system focused primarily on forecasting rainfall for selected cities. Your timestamp is September 30, 2023, meaning you only consider information available prior to this date. I will provide you with the Nino3.4, Pacific Decadal Oscillation (PDO), and North Atlantic Oscillation (NAO) indices for the prediction period {October 1, 2023, to October 15, 2023}. Please integrate this information, consider their climate teleconnection relationship with potential regional rainfall, and combine it with your own knowledge to make a final prediction. For the time being, please ignore narrative responses; I am only interested in numerical results.

—potential forecast—

Period: {October 1, 2023 to October 15, 2023}

Nino3.4: {v1, v2, v3, v4, v5, v6, v7, v8, v9, v10, v11, v12, v13, v14, v15}

PDO: {v1, v2, v3, v4, v5, v6, v7, v8, v9, v10, v11, v12, v13, v14, v15}

NAO: {v1, v2, v3, v4, v5, v6, v7, v8, v9, v10, v11, v12, v13, v14, v15}

Please use the supplied data to predict the rainfall for the above period.

3.5. Baseline comparison: 30-year historical average

To assess the accuracy and trend consistency of GPT's predictions, we used the average daily/monthly rainfall over the past 30 years as a baseline [22]. This standard is widely recognized in climatological studies for defining long-term trends, offering a stable reference point that smooths out short-term anomalies. By comparing GPT's forecasts to this baseline, we aimed to evaluate whether its predictions aligned with climatological norms or exhibit meaningful deviations under different experimental conditions.

We used Root Mean Square Error (RMSE) to measure the error between different predictions and actual rainfall. Additionally, we evaluated the accuracy and trend-capturing ability of the predictions using Pearson's correlation coefficient and Nash-Sutcliffe's efficiency coefficient.

4. Results

Figures 2 and 3 present comparisons across scales and scenarios. In this study, the EM model utilized an LSTM structure, leveraging historical temperature and rainfall data to predict future rainfall. This time series-based deep learning model effectively captured the temporal dependencies within climate data, making it particularly suitable for detecting long-term trends and periodic patterns in meteorological data. As shown, the EM model achieved the best results for both long-term and short-term predictions. We were interested in examining the differences in LLM inferences when professional knowledge was directly incorporated compared to when it was not.

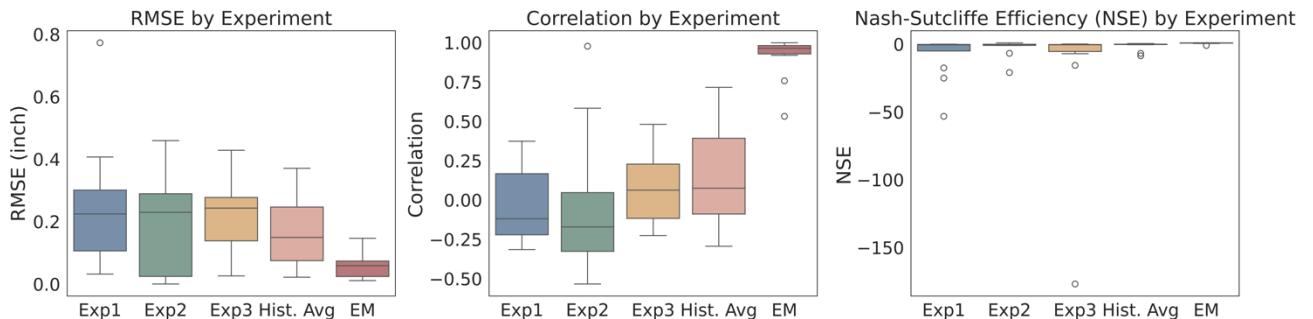


Figure 2. Performance comparison of rainfall prediction across experimental conditions (short-term).

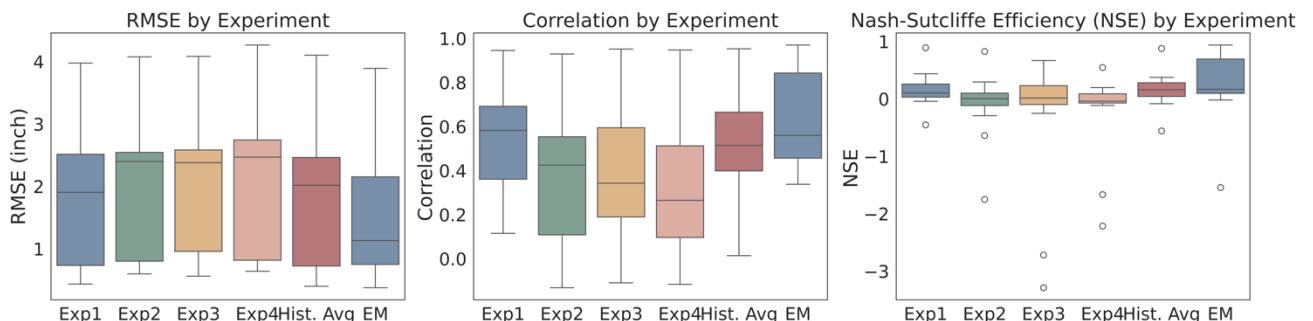


Figure 3. Performance comparison of rainfall prediction across experimental conditions (long-term).

However, when comparing Exp1 and Exp2, we found that providing GPT with relevant domain-specific knowledge and asking it to integrate this with its own knowledge did not significantly improve the results over relying solely on GPT. In short-term predictions, the average RMSEs across all cities for Exp1 and Exp2 were 0.23 and 0.20, respectively, both notably higher than the EM model's 0.06. For long-term predictions, adding domain-specific knowledge increased RMSE compared to Exp1 and reduced the correlation coefficient. Although the provided domain knowledge might offer better predictive insights, results from both short-term and long-term predictions indicated that GPT's internal knowledge plays a dominant role in its inferences.

We were also interested in observing what results GPT could provide when we input factors related to rainfall indirectly than inputting rainfall data directly. Specifically, for long-term predictions, we added teleconnection-related factors (Exp4). In short-term predictions, the indirect provision of regional meteorological factors had limited impact on results; Exp3 showed an RMSE similar to Exp1 and Exp2, with a slightly improved correlation coefficient. However, in long-term predictions, whether adding regional factors or global teleconnection factors, GPT's results declined compared to when this knowledge was not added.

Comparing the results of different experiments with the 30-year historical average, we observed that GPT-generated predictions closely align with the 30-year average in both short-term and long-term forecasts. This similarity suggests that, in the absence of strong trends or notable anomalies, GPT tends to generate conservative predictions that resemble long-term statistical averages. This tendency is even more pronounced in long-term forecasts: We calculated the correlation coefficients between each experiment's results and the 30-year average, finding values of 0.86 for Exp1, 0.82 for Exp2, 0.76 for Exp3, and 0.62 for Exp4, while the EM model showed a lower correlation of only 0.59. This may be due to GPT's limited understanding of physical climate processes, leading it to default to safer, historically consistent forecasts. Additionally, this tendency reflects GPT's low sensitivity to complex climate patterns; in scenarios with limited data or high uncertainty, it relies on historical averages for robustness. However, this approach limits GPT's ability to capture potential trend shifts and the likelihood of extreme climate events.

We further explored the differences between the results generated by GPT-4o and those from the EM model at each time point. In Figures 4 and 5, we compare the time series of short-term and long-term predictions. For the short-term predictions, we found that Exp2 tended to produce results closer to the multi-year average, especially during peak phases in the series. For peak values in the EM predictions, Exp2 noticeably dampened the magnitude of these peaks. For example, in Atlanta, there was an increase in rainfall on October 11 and 12, 2023, relative to the preceding days. The EM model effectively captured this upward trend, while Exp2 significantly reduced the peak values on these dates, bringing them down from the EM's 0.5–0.6 range to the 30-year average level of around 0.1–0.2. Similar patterns were observed on October 11 in Pensacola, October 6 and 14 in New York, and October 4 in Dallas. For peak values that were incorrectly predicted by the EM model, Exp2 also reduced their magnitude, as seen on October 15 in Phoenix and October 4 in Tucson.

We also found that Exp2's tendency to generate results close to the multi-year average was more pronounced in cities with higher rainfall. For instance, in Mobile and Baton Rouge, the results generated by Exp2 at each time point were closer to the 30-year average. A similar pattern could also be observed in the long-term predictions, where for all cities, Exp2's monthly-scale results were closer to the multi-year average at each time point. This tendency likely reflects GPT's preference for a more stable prediction aligned with the historical average than adjusting fully to the extreme values suggested by the professional model. It also indicates that when encountering anomalies or extreme values, GPT is inclined to revert toward the average, leading to a smoothing effect in its predictions.

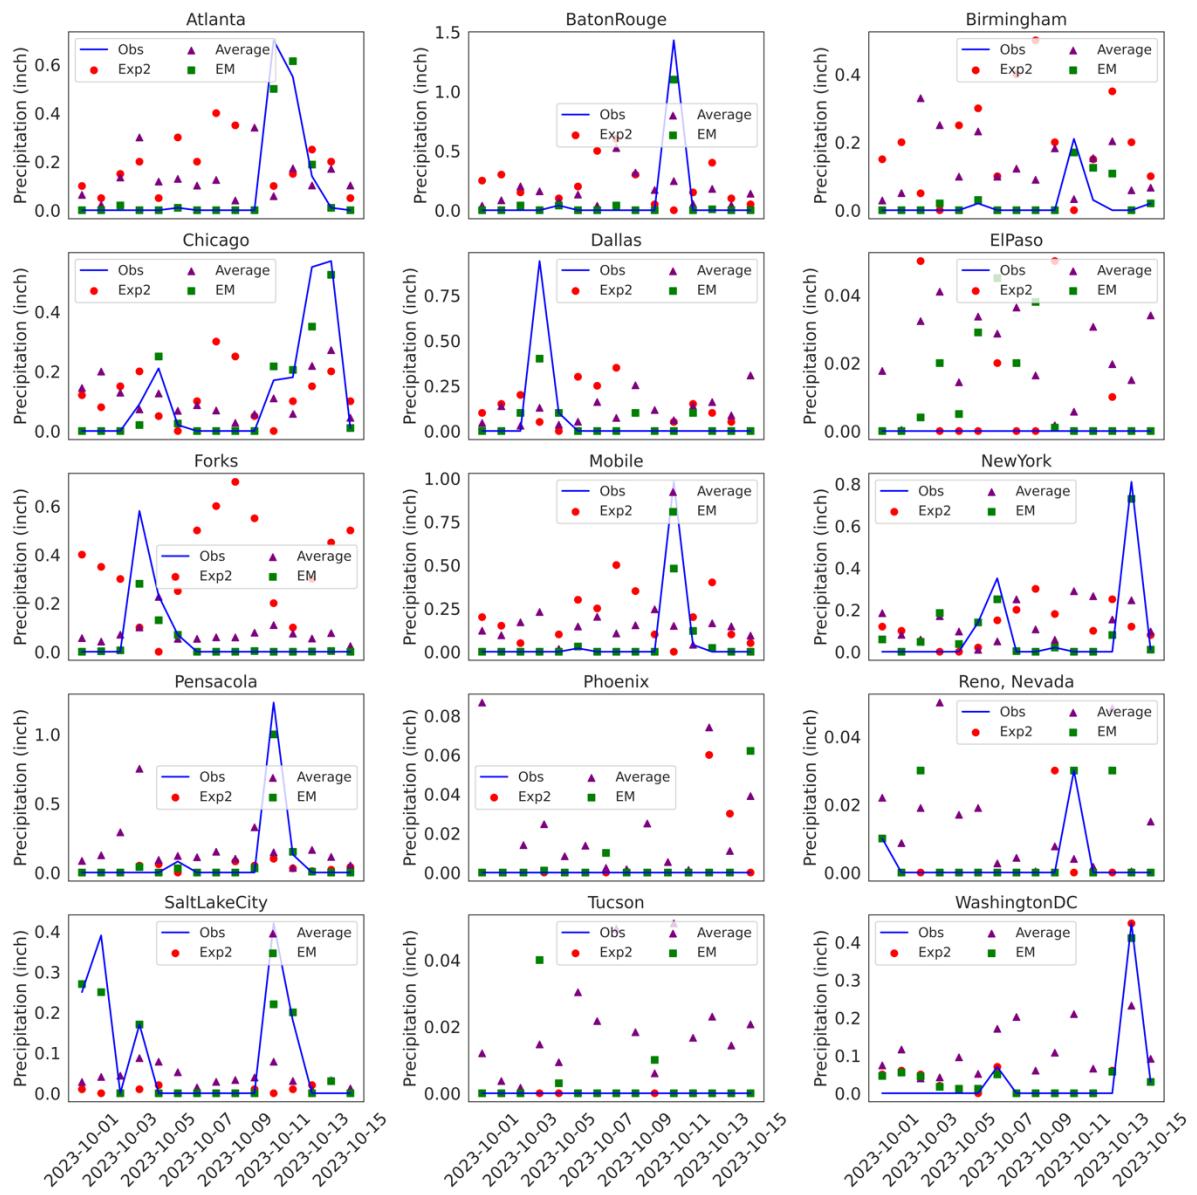


Figure 4. Time series comparison of short-term rainfall predictions across experimental conditions for different cities.

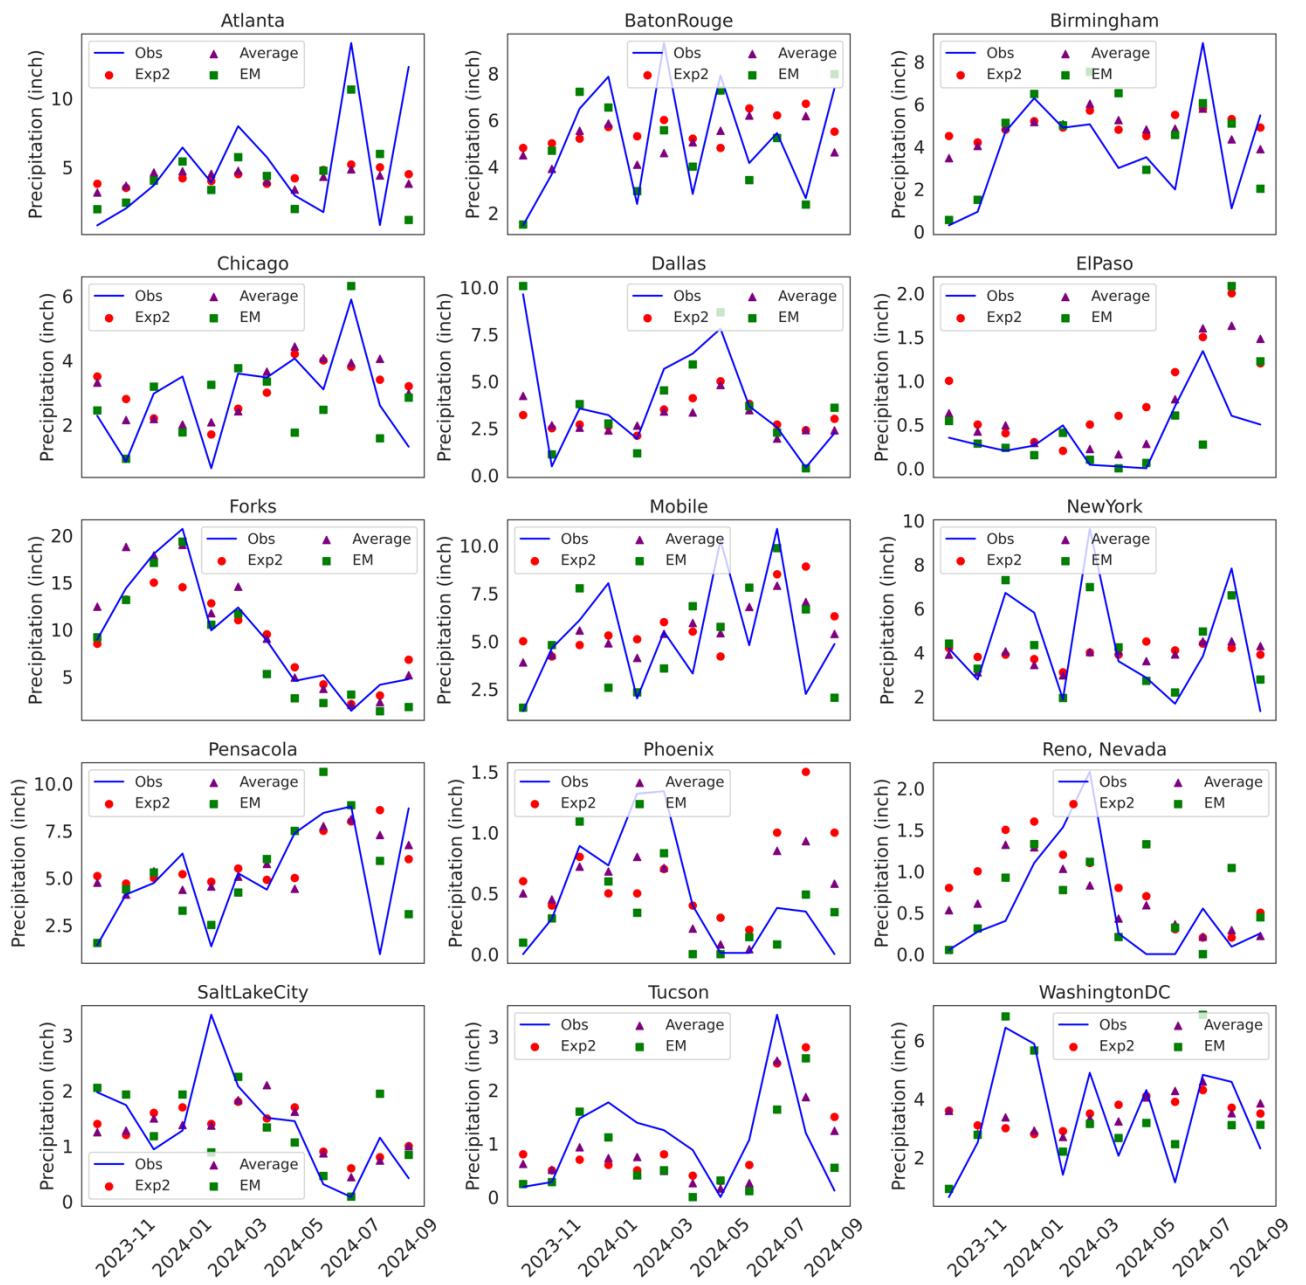


Figure 5. Time series comparison of long-term rainfall predictions across experimental conditions for different cities.

5. Discussion and future research

As a large language model, GPT is fundamentally trained to learn language patterns from vast amounts of textual data, rather than physical processes. This means that it lacks the inherent physical constraints present in climate systems (such as energy conservation, atmospheric and ocean dynamics) and cannot consider causal relationships in predictions as physical models do. Thus, when predicting future rainfall trends, GPT primarily relies on “patterns in text” and “common-sense associations.” When it fails to detect clear trend signals, it defaults to generating results aligned with historical averages—a relatively safe and conservative inference approach that avoids extreme or highly volatile

predictions. This strategy reduces the risk of producing anomalous forecasts and limits GPT's sensitivity to changing trends.

GPT's generation mechanism, based on language patterns, is more adept at producing general and trend-based content, yet often lacks sensitivity to rare or anomalous events. Extreme climate events are relatively infrequent in historical data, so without specialized training, GPT may lean toward predictions near the average, thus minimizing the risk of extreme error. This tendency causes GPT to smooth out results when dealing with extreme rainfall events, missing abnormal signals and downplaying the significance of extreme events. When directly provided with rainfall predictions from expert models (such as the EM model), GPT may not accurately interpret the data or effectively adjust its prediction strategy based on it. In other words, GPT is likely to treat this input as textual information, struggling to extract useful climate patterns or physical principles from it. This limitation diminishes its effectiveness in integrating specialized knowledge into its predictions.

In our study, we experimented with an alternative approach for long-term monthly-scale predictions. We calculated the standard deviation of historical rainfall for each calendar month. A high standard deviation indicates greater rainfall variability for that month, which corresponds to a higher prediction difficulty, while a low standard deviation suggests that rainfall levels are more consistent, indicating potentially lower prediction difficulty [16]. We used this standard deviation as a representation of potential uncertainty and input it into GPT with the EM's predicted rainfall using the following prompt:

Prompt Sample 5:

You are a climate data prediction system focused primarily on forecasting rainfall for selected cities. Your timestamp is September 30, 2023, meaning you only consider information available prior to this date. I will provide you with a potential {daily} prediction for the period {October 1, 2023 to October 15, 2023} based on a deep learning model for the {city}. The standard deviation here can be used as a measure of uncertainty. A smaller standard deviation indicates higher predictability, suggesting that my model's result has lower uncertainty. Conversely, a larger standard deviation indicates greater difficulty in prediction, meaning higher uncertainty in my model's results. Please focus on this measure of uncertainty, and combine it with your knowledge, such as historical trends, to make the final prediction. Please consider the results of the model and combine them with your knowledge to make a final forecast. For the time being, please ignore narrative responses; I am only interested in numerical results.

—potential forecast—

Period: {October 1, 2023 to October 15, 2023}

Rainfall: {v1, v2, v3, v4, v5, v6, v7, v8, v9, v10, v11, v12, v13, v14, v15}

Standard Deviation: {s1, s2, s3, s4, s5, s6, s7, s8, s9, s10, s11, s12, s13, s14, s15}

Please use the supplied data to predict the rainfall for the above period.

This setup provided improved results. As shown in Figure 6, after adding the STD information, GPT's results were close to those of the EM, with a significant improvement compared to Exp2. We present the time series comparison for the three cities with the most significant RMSE improvement in Figure 7. We observed that after adding this uncertainty information, GPT tended to provide results closer to those of the EM. However, we did not achieve the expected outcome where GPT would adjust the predictions with greater potential uncertainty.

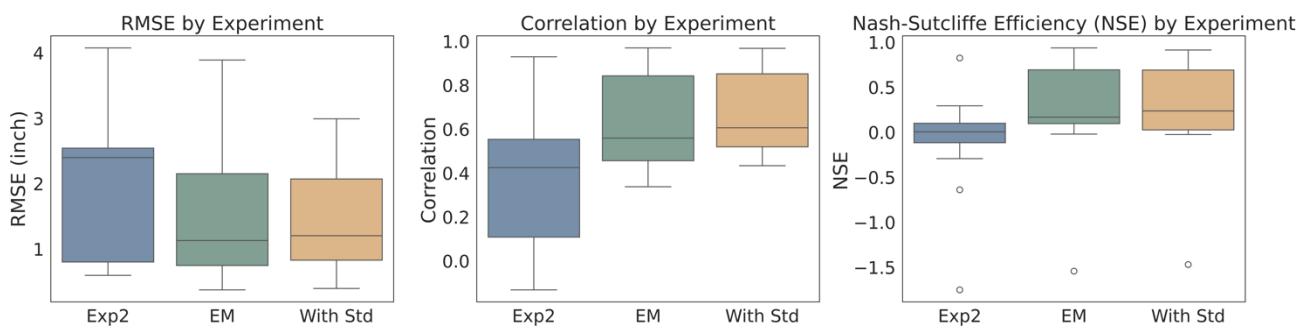


Figure 6. Performance comparison of rainfall prediction across Exp2, EM, and another experiment of adding standard deviation.

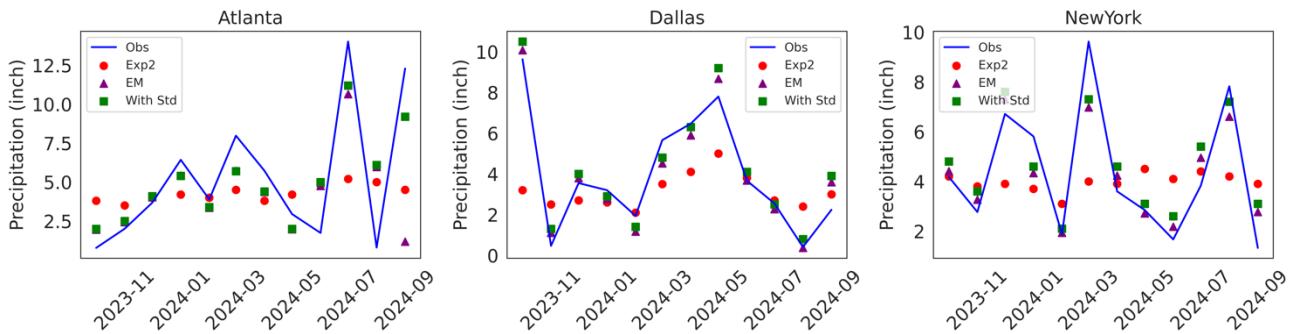


Figure 7. Time Series Comparison of long-term Rainfall Predictions across Exp2, EM, and another experiment of adding standard deviation.

A potential framework to integrate GPT’s knowledge with EM’s predictive insights is to leverage their strengths across uncertainty intervals, resulting in more reliable predictions. By having the EM model provide uncertainty information along with its predictions, we can effectively combine GPT’s conservative prediction tendencies with the high-precision predictions of the physical model. Specifically, the EM model’s predictions can use uncertainty information to define “confidence regions”. In low-uncertainty regions, GPT can directly adopt the EM model’s predictions, as the EM model is generally reliable in these areas. In high-uncertainty regions, however, GPT can use its own knowledge to generate relatively conservative predictions, defaulting to historical averages or stable values to mitigate reliance on extreme fluctuations. By positioning the EM model as the “high-confidence predictor” and GPT as the “smoothing factor” in high-uncertainty scenarios, this strategy balances model accuracy with stability, enhancing overall prediction reliability.

In this study, we analyzed only a limited set of cities and time series, which imposes certain constraints. Additionally, we provided domain knowledge through direct prompts. In the future, optimization strategies such as knowledge distillation, prompt engineering, or multi-task learning could be explored to improve GPT’s understanding and handling of climate data [23]. For instance, incorporating a multi-task learning framework might allow GPT to autonomously learn the weights and importance of climate factors when processing climate data, potentially helping it to better integrate domain-specific knowledge and demonstrate greater potential in climate forecasting.

6. Conclusions

We investigated the ability of large language models (LLMs), specifically ChatGPT-4o, to provide future climate information, focusing on rainfall prediction accuracy. We utilized a 2-layer

LSTM model as an Expert Model (EM), assuming it could generate reliable future predictions. Through a series of experiments, we compared ChatGPT-4o's rainfall predictions under varying conditions: relying solely on its internal knowledge, directly receiving rainfall predictions from the EM, and indirectly inferring rainfall from other related factors predicted by the EM. In these experiments, we evaluated both short-term (15-day, daily scale) and long-term (12-month, monthly scale) prediction capabilities. The results revealed that ChatGPT-4o consistently prioritizes stable predictions closely aligned with historical averages, regardless of whether it integrates additional information from the EM. This tendency highlights the LLM's inherent bias towards conservative outputs, which may limit its effectiveness in capturing dynamic or extreme variations in climate scenarios.

Use of AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools for this article.

Conflict of interest

The authors declare no conflict of interest.

The data were sourced from the following publicly available databases:

- <https://www.ncei.noaa.gov/pub/data/ghcn/daily/>,
- <https://doi.org/10.1184/R1/7890488.v6> [24],
- <https://www.weather.gov/wrh/climate?wfo=sew>.

References

1. S. D. Campbell, F. X. Diebold, Weather forecasting for weather derivatives, *J. Am. Stat. Assoc.*, **100** (2005), 6–16. <https://doi.org/10.1198/016214504000001051>
2. C. Hewitt, S. Mason, D. Walland, The global framework for climate services, *Nature Clim. Change*, **2** (2012), 831–832. <https://doi.org/10.1038/nclimate1745>
3. OpenAI, J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, et al., GPT-4 technical report, arXiv: 2303.08774. <https://doi.org/10.48550/arXiv.2303.08774>
4. J. K. Kim, M. Chua, M. Rickard, A. Lorenzo, ChatGPT and large language model (LLM) chatbots: the current state of acceptability and a proposal for guidelines on utilization in academic medicine, *J. Pediatr. Urol.*, **19** (2023), 598–604. <https://doi.org/10.1016/j.jpurol.2023.05.018>
5. M. Leippold, Thus spoke GPT-3: interviewing a large-language model on climate finance, *Financ. Res. Lett.*, **53** (2023), 103617. <https://doi.org/10.1016/j.frl.2022.103617>
6. A. J. Thirunavukarasu, D. S. J. Ting, K. Elangovan, L. Gutierrez, T. F. Tan, D. S. W. Ting, Large language models in medicine, *Nat. Med.*, **29** (2023), 1930–1940. <https://doi.org/10.1038/s41591-023-02448-8>
7. H. Zhu, P. Tiwari, Climate change from large language models, arXiv: 2312.11985. <https://doi.org/10.48550/arXiv.2312.11985>
8. M. Kraus, J. Bingler, M. Leippold, T. Schimanski, C. Senni, D. Stammbach, et al., Enhancing large language models with climate resources, arXiv: 2304.00116. <https://doi.org/10.48550/arXiv.2304.00116>
9. N. Koldunov, T. Jung, Local climate services for all, courtesy of large language models, *Commun. Earth Environ.*, **5** (2024), 13. <https://doi.org/10.1038/s43247-023-01199-1>
10. D. Thulke, Y. Gao, P. Pelser, R. Brune, R. Jalota, F. Fok, et al., Climategpt: towards AI synthesizing interdisciplinary research on climate change, arXiv: 2401.09646. <https://doi.org/10.48550/arXiv.2401.09646>

11. N. Webersinke, M. Kraus, J. A. Bingler, M. Leippold, Climatebert: a pretrained language model for climate-related text, arXiv: 2110.12010. <https://doi.org/10.48550/arXiv.2110.12010>
12. H. Nguyen, V. Nguyen, S. López-Fierro, S. Ludovise, R. Santagata, Simulating climate change discussion with large language models: considerations for science communication at scale, *Proceedings of the 11th ACM Conference on Learning @ Scale*, 2024, 28–38. <https://doi.org/10.1145/3657604.3662033>
13. S. E. Brownell, J. V Price, L. Steinman, Science communication to the general public: why we need to teach undergraduate and graduate students this skill as part of their formal scientific training, *J. Undergrad. Neurosci. Educ.*, **12** (2013), E6–E10.
14. A. Sherstinsky, Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network, *Physica D*, **404** (2020), 132306. <https://doi.org/10.1016/j.physd.2019.132306>
15. Y. Wang, H. A. Karimi, Impact of spatial distribution information of rainfall in runoff simulation using deep learning method, *Hydrol. Earth Syst. Sci.*, **26** (2022), 2387–2403. <https://doi.org/10.5194/hess-26-2387-2022>
16. M. Zhang, J. D. Rojo-Hernández, L. Yan, Ó. J. Mesa, U. Lall, Hidden tropical pacific sea surface temperature states reveal global predictability for monthly precipitation for sub-season to annual scales, *Geophys. Res. Lett.*, **49** (2022), e2022GL099572. <https://doi.org/10.1029/2022GL099572>
17. R. Krishnan, M. Sugi, Pacific decadal oscillation and variability of the Indian summer monsoon rainfall, *Clim. Dynam.*, **21** (2003), 233–242. <https://doi.org/10.1007/s00382-003-0330-8>
18. R. M. Trigo, J. L. Zêzere, M. L. Rodrigues, I. F. Trigo, The influence of the north atlantic oscillation on rainfall triggering of landslides near Lisbon, *Nat. Hazards*, **36** (2005), 331–354. <https://doi.org/10.1007/s11069-005-1709-0>
19. S. McGregor, C. Cassou, Y. Kosaka, A. S. Phillips, Projected ENSO teleconnection changes in CMIP6, *Geophys. Res. Lett.*, **49** (2022), e2021GL097511. <https://doi.org/10.1029/2021GL097511>
20. K. M. Lau, H. Weng, Recurrent teleconnection patterns linking summertime precipitation variability over East Asia and North America, *J. Meteorol. Soc. Japan*, **80** (2002), 1309–1324. <https://doi.org/10.2151/jmsj.80.1309>
21. S. Y. Wang, L. E. Hipps, R. R. Gillies, X. Jiang, A. L. Moller, Circumglobal teleconnection and early summer rainfall in the US Intermountain West, *Theor. Appl. Climatol.*, **102** (2010), 245–252. <https://doi.org/10.1007/s00704-010-0260-4>
22. J. Wen, C. Wan, Q. Ye, J. Yan, W. Li, Disaster risk reduction, climate change adaptation and their linkages with sustainable development over the past 30 years: a review, *Int. J. Disaster Risk Sci.*, **14** (2023), 1–13. <https://doi.org/10.1007/s13753-023-00472-3>
23. Q. Wu, G. Bansal, J. Zhang, Y. Wu, B. Li, E. Zhu, et al., Autogen: enabling next-gen llm applications via multi-agent conversation framework, arXiv: 2308.08155. <https://doi.org/10.48550/arXiv.2308.08155>
24. Y. Lai, D. A. Dzombak, Use of historical data to assess regional climate change, *J. Climate*, **32** (2019), 4299–4320. <https://doi.org/10.1175/JCLI-D-18-0630.1>



AIMS Press

©2025 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>)