Beyond Accuracy: A Comparative Study of Machine Learning Models for Extreme Weather Forecasting in Rio de Janeiro with a Green AI Perspective

Gabriel B. Breder[1], Gyslla Vasconcelos[1], Mariza Ferro[1]

[1]Instituto de Computação – Universidade Federal Fluminense (UFF)
Av. Gal. Milton Tavares de Souza, S/Nº, São Domingos – Niterói – RJ – Brasil.

{gabrielbertobreder, gysllav, mariza_ferro}@id.uff.br

Abstract. *In recent years, climate change has severely impacted the city of Rio de Janeiro. Some extreme events have occurred, such as temperatures above 40° C and heavy rains that have caused landslides, floods, and deaths. This study investigates the application of Machine Learning (ML) models (MLP, XG-Boost, and LSTM) to predict extreme temperature and precipitation for the city of Rio de Janeiro, while critically evaluating their environmental impact through a Green AI perspective. Beyond traditional accuracy metrics like RMSE, our aim is to assess the trade-off between good prediction and computational efficiency, energy consumption, CO2 emissions, and water usage during model training.*

1. Introduction

Extreme weather events, intensified by climate change, are increasingly responsible for natural disasters that cause significant material damage and loss of life [Intergovernmental Panel on Climate Change (IPCC) 2021]. These events disproportionately affect cities in the global south, such as Rio de Janeiro, Brazil, where rapid urbanization and socio-spatial inequalities exacerbate vulnerabilities [Sharma et al. 2024, Regoto et al. 2021]. In Rio de Janeiro, extreme rainfall during spring and summer has historically led to floods, landslides, and other disasters, particularly in low-income areas located in geologically unstable regions [Dereczynski and Calado 2017]. To address these challenges, the city established the Rio Operation Center (COR [1]), which collaborates with Alerta Rio [2] to provide weather alerts based on telemetry stations, weather radars, and Numerical Weather Prediction (NWP) models.

In recent years, there has been an alarming increase in heat waves that requires new measures in municipal operations. In March 2024, Rio recorded a historic heat index of 62.3°C, the highest since the measurements began in 2014 [3]. By September 2024, winter temperatures reached 39.4°C, the second hottest temperature of the year [4], prompting COR to implement heat-level classifications and alert protocols similar to those adopted for rainfall.

These developments highlight the critical role of data-driven governance in enabling timely public safety measures and mitigating the social impacts of extreme weather.

[1]Centro de Operações Rio - `https://cor.rio/`
[2]Alerta Rio - `http://alertario.rio.rj.gov.br/`
[3]News article about a heat wave in Rio de Janeiro (in Portuguese), available at: link
[4]News article about second hottest temperature of the year, available at: link

Forecast models are an essential component for early warning systems and consecutive actions within crisis management and risk prevention. However, the precision of current forecasting systems for extreme conditions remains insufficient, particularly in complex urban environments such as Rio de Janeiro, where topography, ocean proximity, and microclimates impose even greater challenges to this prediction [da Silva et al. 2022a].

Artificial Intelligence (AI), especially Machine Learning (ML), offers transformative potential for improving meteorological forecasts [Han et al. 2023, Kumar et al. 2024] and supports Sustainable Development Goals (SDG), notably SDG 11 and SDG 13 [Vinuesa et al. 2020]. While ML models can enhance accuracy and require fewer resources than traditional NWP models [Allen et al. 2025], they still demand considerable computational power and energy, challenging their use in developing countries [Lamptey et al. 2024]. Deep Learning (DL) models, in particular, have high energy consumption, contributing to CO2 emissions and raising climate concerns [Cowls et al. 2021, Sætra 2021]. Additionally, AI's substantial freshwater usage for data center cooling poses environmental risks [Li et al. 2025]. These sustainability trade-offs are especially critical for resource-limited regions, perpetuating vulnerability.

The climate impact of these models is controversial in projects addressing climate change. A key factor in choosing them for city governments is the operational cost. These models require retraining for each new extreme event to capture evolving dynamics, which may be unfeasible in developing countries [Lamptey et al. 2024]. This technological gap can perpetuate a cycle where the most affected regions have the least access to advanced forecasts. High computational costs deepen this disparity, leaving vulnerable populations without effective early warning systems.

Therefore, the general objective of this work is to investigate the use of ML models to forecast extreme weather events (temperature and precipitation) in Rio de Janeiro, while critically evaluating their environmental sustainability. Beyond traditional accuracy metrics (e.g., RMSE), the work introduces a Green AI perspective, comparing the computational efficiency, energy consumption, CO2 emissions, and water usage of three ML models (XGBoost, MLP, and LSTM). By addressing the trade-offs between predictive performance and ecological impact, the paper aims to demonstrate that climate resilience solutions can and must balance technical efficacy with environmental responsibility, particularly in resource-constrained regions. This study also highlights the challenges of data imbalance in extreme weather prediction and advocates for equitable access to sustainable ML tools, aligned with SDG 11 and SDG 13.

This work is organized as follows. In Section 2 the related work is presented; Section 3 shows the methodology, data preprocessing, model configuration, and the evaluation used. In Section 4 the experiments and results are presented. Finally, Section 5 presents the final considerations and future works.

2. Related Works

Predicting extreme weather events is particularly challenging in urban areas due to the complexity of atmospheric phenomena and the need for timely alerts [Porto et al. 2022, Sharma et al. 2023]. Accurate forecasts of extreme temperature and precipitation are essential for risk assessment and issuing warnings, given their direct impacts on human life and their social and economic consequences. Mountain regions, for example, are especially vulnerable due to orographic effects, often facing intense rainfall and flooding [Sharma et al. 2023]. In this context, we analyzed studies on extreme event prediction in Rio de Janeiro to identify state-of-the-art methods and their limitations.

Rio de Janeiro presents forecasting challenges due to its complex topography and coastal location, which affect atmospheric dynamics [Porto et al. 2022, Dereczynski and Calado 2017]. Several studies address extreme rainfall and temperature in the city. [da Cunha et al. 2024] analyzed trends in extreme climate indices. [da Silva et al. 2022b] evaluated the WRF model, highlighting microphysics and atmospheric mechanisms. [Dereczynski and Calado 2017] emphasized the role of circulation and relief.[Polifke et al. 2020] and [da Silva et al. 2020] explored temperature variability and radar-based approaches to improve rainfall forecasting.

Due to the limitations of traditional methods, there is growing interest in applying AI, especially ML and DL, to forecast extreme weather events using large meteorological datasets [Porto et al. 2022, Bouallègue et al. 2024]. ML models can learn non-linear relationships between variables, offering faster and often more accurate forecasts. However, the rarity of extreme events poses data scarcity challenges, addressed through approaches like meta-learning and few-shot learning [Bendre et al. 2020]. DL models, such as RNNs, have been used to capture spatio-temporal patterns in rainfall prediction [Shi et al. 2015]. YConvLSTM [Porto et al. 2022] combines outputs from NWP models to improve heavy rainfall detection. Other models, like MetNet [Sonderby et al. 2020] and RainNet [Ayzel et al. 2020], use radar and satellite data with advanced architectures, yet forecasting rare events remains a challenge requiring further advances.

Therefore, the studies analyzed demonstrate the progress of ML techniques in predicting extreme events, applied to forecasting precipitation and temperature from meteorological data [Sonderby et al. 2020, Ayzel et al. 2020, Porto et al. 2022]. However, a recurring challenge in the literature is data imbalance, which affects the detection of rare events, in addition to the high computational cost of some DL-based models [Bouallègue et al. 2024, Bendre et al. 2020].

The literature also proposed works related to the three models used in this study. For temperature prediction, [Xu et al. 2024] present an LSTM model with high accuracy. [Ma et al. 2020] applied the XGBoost model to predict temperature 1, 2 and 3 hours ahead, also obtaining good accuracy. Furthermore, [Mfetoum et al. 2024] showed an MLP model that had good performance in predicting solar irradiation. Meanwhile, for precipitation, [Anwar et al. 2021] presented an XGBoost model and had faced a really challenging task with many 0 mm precipitation values in the data, such as ours. However, none of the studies showed energy efficiency or time wasted training the models, highlighting a gap that this work aims to explore.

In this work, we explore three ML models to predict temperature and precipitation in the city of Rio de Janeiro. Beyond evaluating the accuracy, as in related work, we evaluated the computational and environmental impact of the approaches. In this way, we seek to contribute to the literature by comparing different techniques, considering predictive efficiency and sustainability.

3. Methodology

The methodology proposed to achieve the objectives of this work begins with Data Exploration and Analysis (3.1), followed by Data Preprocessing (3.1); Modeling for three ML models trained and tested (3.2) which are quantitatively evaluated - Model Evaluation (3.3). This pipeline was designed to handle the challenges of weather prediction using ML models. The problem addressed here consists of a supervised ML task, more specifically regression, since we want to predict the value of precipitation in millimeters (mm) and the temperature in degrees. The study area is the city of Rio de Janeiro, and for this work is the neighborhood of Copacabana.

3.1. Data Exploration and Preprocessing

The experiments involved the use of data from meteorological stations of the National Institute of Meteorology (INMET) [5] downloaded by year, and a script was created to join the files of the Forte de Copacabana station, Rio de Janeiro. For this work, the Copacabana subset was selected in the period of May 18, 2007 to February 16, 2024 with the target features: **Precipitation** and **Temperature**, which were trained separately. This period was selected due to its low number of missing data. After this period, from March 19, 2024 to July 30, 2024 there is a huge amount of missing data with no values representing radiation, temperature, and humidity features.

This station was selected among other four INMET stations, because it is located in a coastal region, close to some other sensors, such as the meteorological buoys in Guanabara Bay [6] and the pluviometric stations of AlertaRio. Furthermore, the Copacabana station contains data collected from 2007 onward, that is, just no more than Marambaia with data from 2002. However, the Marambaia station is not well located when considering the distance to those other stations. In the future, in addition to the weather stations, we plan to integrate these complementary data sources to train ML models for extreme precipitation forecast. This concern with the location of the stations came due to the challenge of training ML models with spatially non-uniformity observations collected in regions of space that do not form a regular grid.

This database includes 19 features and 145,139 instances measured by meteorological instruments with a temporal resolution of 1 hour, characterizing a time series dataset. Table 1 provides details on the features and their descriptions. Analyzing the number of missing values for each feature in this subset of Copacabana, only the Radiation feature had a significant amount of missing data - 36% of all data. In a deep analysis, it was observed that most of these missing data were about the period of night when there is no sun, so the missing values were set to 0. The other features had only a small amount of missing data with less than 5% each and were filled using the **linear spline interpolation**.

The following analysis was about the distribution of the precipitation data since one of the main focus of the study was to predict extreme rainfall. According to AlertaRio's classification criteria [7], precipitation accumulated in 1 hour is categorized into

[5] INMET - `https://portal.inmet.gov.br/dadoshistoricos`

[6] These sensors collect water temperature, among others, which were indicated by the project meteorologists as relevant for predicting extreme events in Rio de Janeiro.

[7] Preciptation Distribution - `https://www.sistema-alerta-rio.com.br/`

Table 1. Meteorological data from INMET stations – Features and description

#	Feature	Description
1	Date	Date of the measurement
2	Time	Time (UTC) of the measurement
3	Precipitation	Hourly precipitation (mm)
4	Pressure	Atmospheric pressure at station level (mB)
5	MaxPressure	Max atmospheric pressure in the last hour (mB)
6	MinPressure	Min atmospheric pressure in the last hour (mB)
7	Temperature	Air temperature (°C)
8	MaxTemp	Max air temperature in the last hour (°C)
9	MinTemp	Min air temperature in the last hour (°C)
10	DewPoint	Dew point temperature (°C)
11	MaxDewPoint	Max dew point in the last hour (°C)
12	MinDewPoint	Min dew point in the last hour (°C)
13	Humidity	Relative humidity (%)
14	MaxHumidity	Max relative humidity in the last hour (%)
15	MinHumidity	Min relative humidity in the last hour (%)
16	WindSpeed	Hourly wind speed (m/s)
17	WindDirection	Hourly wind direction (°)
18	MaxWindGust	Max wind gust in the last hour (m/s)
19	Radiation	Global radiation (KJ/m²)

Figure 1. Temperature vs. Humidity

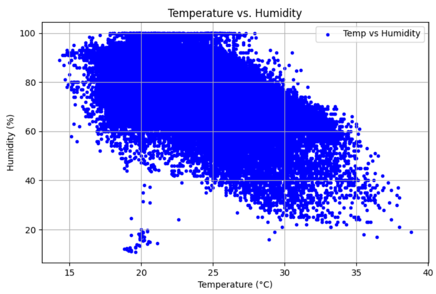

five groups, as shown in Table 2. This classification was utilized in order to exhibit how difficult it is to make a regression model that can predict moderate, heavy, or violent rain with a limited amount of data.

Table 2. Rain rate statistics of data - Forte de Copacabana meteorological station - monitored by INMET from May 18, 2007 to February 16, 2024.

Rain Rate (mm/h)	Proportion (%)	Severity Level
$x = 0$	91.42	No Rain
$0 < x < 5$	8.01	Light Rain
$5 \leq x < 25$	0.53	Moderate
$25 \leq x < 50$	0.03	Heavy
$x \geq 50$	0.01	Violent

The data distribution is highly uneven with 91% of the days without any precipitation (Table 2). This significant data imbalance is a huge challenge for ML models. On the other hand, Figure 1 illustrates that temperature and humidity are strongly correlated, where a decrease in temperature affects an increase in humidity. This correlation suggests that predicting the temperature may be easier than predicting the precipitation.

The features of the database and its values changed after 2019, which required

previsao-do-tempo/termosmet/

standardization. The name of column “Hour (UTC)” changed to “Hour UTC” and the hour value represented by “hh:mm” changed to “hhmm UTC”. Additionally, float values that were represented by “.” turned into “,” and the missing values that before 2019 were represented by “-9999”, began to be represented by a blank space. As the database splits date and time into two different features, it was necessary to perform a **feature engineering** joining them into one called Timestamp.

3.2. Setup of Machine Learning Models

The eXtreme Gradient Boosting (XGBoost) is an implementation of gradient boosted tree algorithms. This is an ensemble learning technique that makes use of a set of base learners to improve the stability and effectiveness of an ML model. It is called gradient boosting because it uses a gradient descent algorithm to minimize loss when adding new models. The central idea of boosting is the implementation of homogeneous ML algorithms in a sequential way, where each tries to improve the model by focusing on the errors made by the previous one [Chen and Guestrin 2016]. The hyperparameters used in the training are detailed in Table 3 defined by means of the Grid Search method.

Table 3. XGBoost Hyperparameters used in the Experiments to predict precipitation (Value Prec) and temperature (Value Temp)

Hyperparameter	Value Prec	Value Temp	Description
eta	0.01	0.1	Learning rate
colsample_bytree	1.0	0.6	Fraction of features used per tree
max_depth	8	6	Maximum tree depth
min_child_weight	5	3	Minimum weight sum in a child node
subsample	0.6	0.6	Fraction of training samples per tree

Multi-Layer Perceptron (MLP) is a feedforward neural network consisting of three layers: input layer, hidden layer, and output layer [Rosenblatt 1958]. For the input and hidden layers the ReLu activation function was utilized and for the output layer, as it is a regression model, the linear activation function was used with one neuron. The hidden layer consists of 32 neurons.

LSTM is neural network used to maintain sequential information; this enables predictions based on past-sequences within a defined window size. The model architecture consisted of two layers, each with a hidden dimension of 256 units, and was trained for 25 epochs. For the temperature prediction model, a learning rate of 0.001 and a batch size of 64 were used. In contrast, the precipitation prediction model employed a learning rate of 0.01 and a batch size of 100.

Models based on decision trees, such as XGBoost, offer greater energy efficiency compared to those based on neural networks. Some studies have shown reasonable results in weather prediction, as presented in Section 2 and are fast model for training. However, neural network-based models have proven to be more precise in handling time series data, despite their higher training costs. Although MLP can manage time series data, they are not the most efficient or natural choice compared to LSTM, MLPs can perform reasonable short-term dependencies like predicting temperature and precipitation a few hours ahead, which aligns with our experiments, and they train faster than LSTM, which is superior for predicting extreme events. The varied complexity and training times led to an evaluation of these three models to balance predictive accuracy with environmental sustainability.

These three ML models were used to predict temperature and precipitation. The dataset was divided into 80% for training and 20% for test, so the data set of May 18, 2007 to October 25, 2020 was used for training and October 24, 2020 to February 16, 2024 for the test phase.

3.3. Quantitative Evaluation

The **Root Mean Square Error (RMSE)** was used to assess the predictive accuracy of the ML models, a common metric for regression tasks. To track energy use and carbon footprint, we employed the Python library `codeCarbon`[8], which monitors emissions with minimal code changes by wrapping model functions. After each run, it reports total energy consumption in kWh. CO_2 equivalent emissions were calculated by multiplying energy use by the local grid's carbon intensity, adjusted by the Power Usage Effectiveness (PUE), which accounts for total infrastructure energy. The global average PUE is published annually by the Uptime Institute [report 2024].

Water consumption is a recent topic and challenging to estimate. To calculate the hydric footprint, we used the equation $w = e \cdot [\rho_{s1} + \theta \cdot \rho_{s2}]$ from [Li et al. 2025], where w is the water footprint, e the energy consumed; ρ_{s1} and ρ_{s2} the on-site and off-site WUE; and θ the PUE.

As experiments were conducted on a local machine in Brazil without cooling systems, the on-site WUE = 0 and PUE = 1. Since Brazil's energy grid is mainly hydroelectric, the off-site WUE = 18.59 [Reig et al. 2020], simplifying the equation to $w = e \cdot \rho_{s2}$. The local hardware used was an **Intel(R) Core(TM) i7-1065G7 CPU @ 1.30GHz** with **12GB** of RAM. The preprocessed dataset and ML code can be found in Github.

4. Experiments and Results

In this section, we report initial experiments with the three ML and setup described in Section 3.2. These experiments aim to highlight the potential of the different models in tackling the challenges in the forecast of extreme weather events (temperature and rainfall) and their energy consumption aligned with carbon emissions and water consumption for training.

Regarding the forecast temperature, as illustrated in Figure 2, the comparison among the three ML models shows that XGBoost and LSTM produced comparable results, while MLP performed significantly worse. The MLP model struggled to accurately predict temperatures below 20° C and above 30° C.

Table 4 provides a comparison of all results in different quantitative metrics. Clearly, the LSTM model outperformed the others in terms of RMSE. However, its execution time was considerably longer, resulting in energy consumption that was 44 times greater than that of XGBoost. While the MLP model delivered inferior results compared to the other two models, it consumed less energy than the LSTM.

In forecasting precipitation, the models struggled to predict precipitation above 5 mm of cumulated rain in one hour. This was caused by the significant imbalance in the data, as illustrated in Table 2, where non-rainfall values constitute over 91% of the total.

[8] `http://codecarbon.io/`

Figure 2. Residual Plot of Temperature

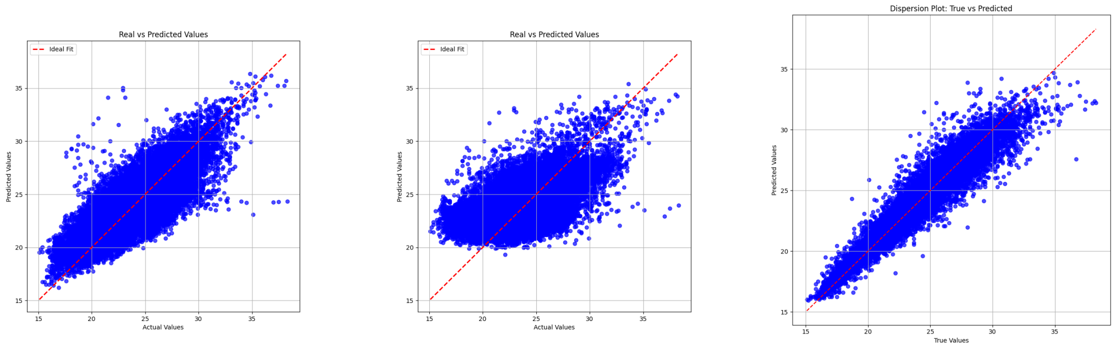

(a) XGBoost **(b) MLP** **(c) LSTM**

Table 4. RMSE, energy consumed in kWh (Energy), CO_2e emission in Kg (CO_2e) and water consumption in Liters (Water) for temperature prediction.

Model	Time(s)	RMSE	Energy(kWh)	CO_2(Kg)	Water(L)
XGBoost	95	1.80	0.000441	0.000043	0.008196
MLP	384	2.46	0.001787	0.000176	0.033202
LSTM	2788	0.80	0.019453	0.001913	0.361631

Table 5. Comparison of RMSE, time in seconds, energy consumed in kWh (Energy), CO_2e emission in Kg (CO_2e) and water consumption in Liters (Water) for precipitation prediction.

Model	Time (s)	RMSE	Energy (kWh)	CO_2 (Kg)	Water (L)
XGBoost	100	1.10	0.000467	0.000046	0.008679
MLP	327	1.15	0.001523	0.000150	0.028313
LSTM	2958	1.12	0.010987	0.001861	0.204248

Comparison of results (Table 5) shows that all models had a really low RMSE error, clearly demonstrating that the models overfitted the data without rain. Despite the low RMSE, the results are not good, as we need to predict when and how much rain will fall and not when the precipitation is closer to 0 mm. We performed a new set of experiments that removed all precipitation values equal to 0 mm. The hypothesis was that by balancing the data, the ML models could improve the learning. However, we got similar results.

In summary, analyzing the data was a really difficult task, with numerous missing data and inconsistencies in sensor data representation across different years. Since temperature has a strong correlation with other characteristics, such as humidity, it was easier to predict. However, rainfall was a big challenge faced, especially because it has a majority of no rain values.

Figure 3 shows a comparison for the three models; the blue bars represent the RMSE error and the orange bars the energy consumed; both bars should be minimized. It's evident that the orange bar representing LSTM in both experiments are the preponderant one, indicating the high energy consumed. As expected, XGBoost consumes less energy among the three models while achieving an accuracy similar to LSTM.

Figure 3. Performance Comparison between XGBoost and LSTM Model.

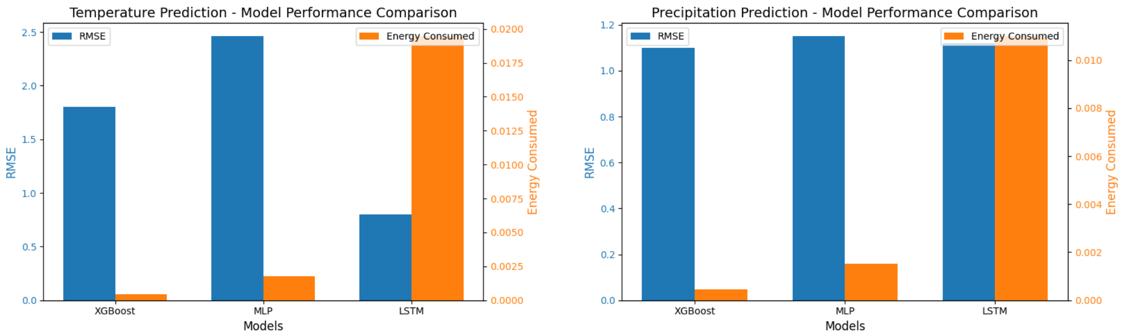

(a) Temperature **(b) Precipitation**

5. Final Considerations and Future Works

This study explored the use of ML models to predict temperature and precipitation in Rio de Janeiro, focusing not only on predictive accuracy but also on the environmental and operational costs of these models. Beyond traditional accuracy metrics (e.g., RMSE), the work advocates for a green AI perspective, comparing computational efficiency, energy consumption, CO2 emissions, and water use. Our aim is to raise the critical need to balance predictive performance with sustainability in climate-related ML applications.

We evaluated three ML models—XGBoost, MLP, and LSTM—for forecasting extreme temperature and precipitation in Rio de Janeiro. For temperature prediction, LSTM achieved the lowest RMSE (0.80), outperforming XGBoost (1.80) and MLP (2.46), demonstrating its strength in capturing temporal patterns. However, for precipitation prediction, all models struggled due to severe data imbalance (91% of the samples had no rain), producing misleadingly low RMSE values. Although these scores suggest high accuracy, they primarily reflect the model's ability to predict "no rain" events, masking their poor performance in extreme rainfall. Regarding the **trade-off between precision and sustainability**, LSTM, despite being more accurate, consumed 44 times more energy than XGBoost, emitting 0.0019 kg CO_2 versus 0.000043 kg, and using 0.36 L of water compared to 0.008 L.

Indeed, while investigating approaches to extreme rainfall events in the city of Rio de Janeiro, we are confronted with a long time series of data, but at the same time with the scarcity of samples of extreme rainfall, resulting in high imbalance data for training. In addition, this is an initial work and only analyses the Copacabana station, out of three others. Forecasting extreme rainfall is a really hard task and is already known to be a challenge, especially when using models that do not inherently use past-sequence data, like XGBoost. Nevertheless, this approach was essential to evaluate not only predictive accuracy but also environmental costs. However, it is evident that the XGBoost model effectively predicted the temperature well and continued to be greener than the others. With the results, even though the LSTM model is slightly more accurate than XGBoost, it requires significantly more energy.

Although this work presents major limitations with respect to the limited dataset, and the low predictive accuracy for precipitation, these initial results still bring important contributions and achieve the objective proposed in this article of discussing the paradoxical situation of using ML to address climate change challenges without discussing the environmental and economic costs involved. It is necessary to consider that despite the numerous benefits that can be achieved with a high-precision solution - something totally desirable - saving lives and local economies, these costs could make the development and operational execution unfeasible in developing countries. This technological inequality can perpetuate a vicious cycle in which less-advantaged regions, which suffer the most from the impacts of climate change, have the least access to advanced forecasting solutions, maintaining the disparity, leaving resource-constrained regions unable to implement effective early warning systems. Consequently, the populations most vulnerable to climate impacts remain marginalized. Furthermore, this work contributes to the debate on the need for projects to prioritize the choice of greener models, whenever possible, and when not, to commit to evaluating the possibility and, above all, making these costs transparent.

In summary, this work demonstrates that it is possible to reconcile environmental responsibility with technical performance, offering a viable and ethical path for the application of AI in operational meteorology. In future work, the aim is to increase the accuracy and performance of the precipitation model, using methods to oversample the data and to have a more balanced precipitation feature. In order to do that, an integration is going to happen with meteorological buoys from SimCosta and pluviometric stations from Alerta Rio. Moreover, related features will be converted into new ones through Feature Construction process, helped by meteorology specialists.

Ethical Declaration: We used the Writefull tool, a large language model based on GPT, integrated with Overleaf to assist in the grammatical correction and improvement suggestions of this text. Each AI suggestion was made based on our own text and was carefully reviewed, which makes us fully responsible for the form and content of the text of this work. In this work, for the training and tuning of the ML models, we consumed 0.0263 kWh of energy, impacted in 0.0025 CO2 emissions and 0.4888 liters of water.

Acknowledgments

This work was supported by the Serrapilheira Institute (Grant No. Serra – 2211-41897), SERRAPILHEIRA/FAPERJ - 2023 (Grant No. E-26/210.242/2024), FAPERJ - Basic Research Support Grant - APQ1 E_13/2023 (Grant No. 210.723/20240), and the Institutional Program for Scientific Initiation Scholarships at UFF (PIBIC-CNPQ).

References

Allen, A., Markou, S., Tebbutt, W., Requeima, J., Bruinsma, W. P., Andersson, T. R., Herzog, M., Lane, N. D., Chantry, M., Hosking, J. S., and Turner, R. E. (2025). End-to-end data-driven weather prediction. *Nature Communications*.

Anwar, M. T., Winarno, E., Hadikurniawati, W., and Novita, M. (2021). Rainfall prediction using extreme gradient boosting. *Journal of Physics: Conference Series*, 1869(1):012078.

Ayzel, G., Scheffer, T., and Heistermann, M. (2020). Rainnet v1.0: A convolutional neural network for radar-based precipitation nowcasting. *Geoscientific Model Development*, 13:2631–2644.

Bendre, N., Marn, H. T., and Najafirad, P. (2020). Learning from few samples: A survey. *arXiv*.

Bouallègue, Z. B. et al. (2024). The rise of data-driven weather forecasting: A first statistical assessment of machine learning–based weather forecasts in an operational-like context. *Bulletin of the American Meteorological Society*, 105(6):E864–E883.

Chen, T. and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16. ACM.

Cowls, J., Tsamados, A., Taddeo, M., and Floridi, L. (2021). The ai gambit: leveraging artificial intelligence to combat climate change—opportunities, challenges, and recommendations. *Ai & Society*, pages 1 – 25.

da Cunha, F. P., Manso, M. C., and Júnior, A. A. O. (2024). Analysis of the occurrence of extreme events in the city of rio de janeiro/rj. *Revista de Gestão Social e Ambiental*, 18(4):e04730.

da Silva, F. P., da Silva, A. S., and da Silva, M. G. A. J. (2022a). Extreme rainfall events in the rio de janeiro city (brazil): description and a numerical sensitivity case study. *Meteorology and Atmospheric Physics*, 134:1–20.

da Silva, F. P., da Silva, A. S., and da Silva, M. G. A. J. (2022b). Extreme rainfall events in the rio de janeiro city (brazil): description and a numerical sensitivity case study. *Meteorology and Atmospheric Physics*, 134(77).

da Silva, F. P., da Silva, M. G. A. J., Filho, O. C. R., et al. (2020). Pattern evaluation of operational radar reflectivity–rainfall relationship for quitandinha river flooding events: Petrópolis, rio de janeiro (brazil). *Environmental Earth Sciences*, 79:525.

Dereczynski, C. P. and Calado, Renata Novaes e Bodstein de Barros, A. (2017). Chuvas extremas no município do rio de janeiro: Histórico a partir do século xix. *Anuário do Instituto de Geociências - UFRJ*, 40(2):17–30.

Han, D., Im, J., Shin, Y., and Lee, J. (2023). Key factors for quantitative precipitation nowcasting using ground weather radar data based on deep learning. *Geoscientific Model Development*, 16(20):5895–5914.

Intergovernmental Panel on Climate Change (IPCC) (2021). *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge.

Kumar, R., Goel, R., Sidana, N., Sharma, A., ghai, S., Singh, T., singh, R., Priyadarshi, N., Twala, B., and Ahmad, V. (2024). Enhancing climate forecasting with ai: Current state and future prospect. *F1000Research*, 13(1094).

Lamptey, B., Sahabi ABED, S., Gudoshava, M., Mutemi, J., Bopape, M.-J., Adefisan, E., Igri, P., Sanda, S., Ndiaye, O., Parker, D., Dougill, A., Fink, A., Knippertz, P., Woolnough, S., and Kolstad, E. (2024). Challenges and ways forward for sustainable weather and climate services in africa. *Nature Communications*, 15:4.

Li, P., Yang, J., Islam, M. A., and Ren, S. (2025). Making ai less ”thirsty”: Uncovering and addressing the secret water footprint of ai models.

Ma, X., Fang, C., and Ji, J. (2020). Prediction of outdoor air temperature and humidity using xgboost. *IOP Conference Series: Earth and Environmental Science*, 427:012013.

Mfetoum, I. M., Ngoh, S. K., Molu, R. J. J., Nde Kenfack, B. F., Onguene, R., Naoussi, S. R. D., Tamba, J. G., Bajaj, M., and Berhanu, M. (2024). A multilayer perceptron neural network approach for optimizing solar irradiance forecasting in central africa with meteorological insights. *Scientific Reports*, 14(1):3572.

Polifke, F., Filho, O. C. R., da Silva, M. G. A. J., et al. (2020). Observed and estimated atmospheric thermodynamic instability using radiosonde observations over the city of rio de janeiro, brazil. *Meteorology and Atmospheric Physics*, 132:297–314.

Porto, F., Ferro, M., Ogasawara, E., Moeda, T., de Barros, C. D. T., Silva, A. C., Zorrilla, R., Pereira, R. S., Castro, R. N., Silva, J. V., Salles, R., Fonseca, A. J., Hermsdorff, J., Magalhães, M., Sá, V., Simões, A. A., Cardoso, C., and Bezerra, E. (2022). Machine learning approaches to extreme weather events forecast in urban areas: Challenges and initial results. *Supercomputing Frontiers and Innovations*, 9(1):49–73.

Regoto, P., Dereczynski, C., Chou, S. C., and Bazzanela, A. C. (2021). Observed changes in air temperature and precipitation extremes over brazil. *International Journal of Climatology*, 41(11):5125–5142.

Reig, P., Luo, T., Christensen, E., and Sinistore, J. C. (2020). Guidance for calculating water use embedded in purchased electricity.

report, U. I. I. (2024). Uptime institute global data center survey results 2024.

Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65 6:386–408.

Sharma, H., Hegde, N., Banerjee, I., Bhattacharyya, A., and Srivastava, P. (2024). Climate-induced migration in the global south: an in-depth analysis. *npj Urban Sustainability*, 4:133.

Sharma, O., Trivedi, D., Pattnaik, S., Hazra, V., and Puhan, N. B. (2023). Improvement in district scale heavy rainfall prediction over complex terrain of north east india using deep learning. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–8.

Shi, X., Chen, Z., Wang, H., et al. (2015). Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1*, pages 802–810, Cambridge, MA, USA. MIT Press.

Sonderby, C. K., Espeholt, L., Heek, J., Dehghani, M., Oliver, A., Salimans, T., Agrawal, S., Hickey, J., and Kalchbrenner, N. (2020). Metnet: A neural weather model for precipitation forecasting. *arXiv preprint arXiv:2003.12140*. Accessed: 6 Feb. 2025.

Sætra, H. S. (2021). Ai in context and the sustainable development goals: Factoring in the unsustainability of the sociotechnical system. *Sustainability*, 13(4).

Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S., Tegmark, M., and Nerini, F. F. (2020). The role of AI in achieving the Sustainable Development Goals. *Nature Communications*, 11(233).

Xu, J., Wang, Z., Li, X., Li, Z., and Li, Z. (2024). Prediction of daily climate using long Short-Term memory (LSTM) model. *Int. J. Innov. Sci. Res. Technol.*, pages 83–90.