

MAJOR PROJECT REPORT

Enhancing Weather Forecasting with LSTM and Random Forest Algorithms



FOUR- YEAR UNDERGRADUTE PROGRAMM

(DESIGN YOUR DEGREE)

SEMESTER – 2

SUBMITTED TO

UNIVERSITY OF JAMMU, JAMMU

BY

NAME OF THE STUDENTS:	SEMESTER:	ROLL NO:
MALHAR KHADYAL	2ND	DYD-23-10
SHUBHAM SHARMA	2ND	DYD-23-19

UNDER THE MENTORSHIP OF

Prof. K.S. Charak

Professor

Dept. of Mathematics

University of Jammu, Jammu

Dr. Jatinder Manhas

Associate Professor

Computer Science and IT

University of Jammu, Jammu

Dr. Sandeep Arya

Sr. Assistant Professor

Dept. of Physics

University of Jammu, Jammu

Submitted on: _____ September, 2024

Certificate

The report titled “*Enhancing Weather Forecasting with LSTM and Random Forest Algorithms*” submitted by the students of Semester II for partial fulfillment of the requirements for the Design Your Degree, Four Year Undergraduate Program running under the ambit of Skill Incubation Innovation Entrepreneurship Development Centre (SIIEDC) at the University of Jammu.

The project is a key component of Semester 2 of the academic program that covers three major domains that is Coding through Gpt 4, Decoding the world through AI and Language to understand nature under the guidance of Prof. K.S Charak, Dr. Jatinder Manhas, Dr. Sandeep Arya respectively. The project report is original and has not been submitted elsewhere for any academic recognition or any other purpose.

Signature of Students:

Malhar Khadyal

Shubham Sharma

Signature of Mentors:

Prof. K.S Charak

Dr. Jatinder Manhas

Dr. Sandeep Arya

Prof. Alka Sharma

Director, SIIEDC, University of Jammu

Acknowledgment

We take this opportunity to thank everyone who has supported and mentored us as we have worked. First and foremost, we sincerely express our gratitude to Almighty God for inspiring us and giving us the inner fortitude needed to overcome this obstacle.

We lack the words to adequately convey our profound gratitude and sincere respect for our mentor, Dr. Jatinder Manhas. His wise counsel, heartfelt support, and knowledgeable counsel were crucial to our effort. Dr. Jatinder Manhas gave us consistent support and motivation despite his hectic schedule. It is a great blessing and an honour for us to have worked under his knowledgeable tutelage.

Prof. Alka Sharma, who is the director of SIIEDC (Design Your Degree) at the University of Jammu in Jammu and a member of the advisory group, is also greatly appreciated. Her insightful counsel, kind assistance, and departmental supply of essential facilities were crucial to this project's success.

We sincerely thank all of the esteemed mentors for their invaluable assistance in carrying out this significant endeavour. We sincerely thank our friends for their collaboration and support; their wonderful companionship, warmth, and love made this study project enjoyable and rewarding.

We thank everyone who has supported us in reaching our objectives, whether they are included here or not, and who has silently sent us best wishes and insightful advice.

Name of group members:

Malhar Khadyal (DYD-23-10)

Shubham Sharma (DYD-23-19)

Abstract

Enhancing Weather Forecasting with LSTM and Random Forest Algorithms

Today's society depends heavily on weather forecasting, which affects industries including transportation, disaster relief, and agriculture. Even with major advances in meteorological research, forecasting weather patterns—particularly rainfall—remains a difficult task. The intrinsic complexity and non-linearity of weather systems may prove to be a challenge for traditional forecasting techniques, which frequently rely on statistical models and physical simulations. By using cutting-edge machine learning methods, this research aims to improve the precision and dependability of weather forecasts.

This project's main goal is to enhance weather forecasting by the application of Random Forest and Long Short-Term Memory (LSTM) networks, two advanced machine learning methods. Recurrent neural networks (RNNs), such as LSTM networks, operate very well with sequential data, which makes them a good choice for time series forecasting. Their purpose is to record temporal patterns and long-term dependencies in data, which is essential for precise weather forecasting. However, Random Forest—an ensemble learning technique based on decision trees—performs exceptionally well in managing high-dimensional data and intricate feature interactions, offering modelling flexibility and robustness.

To achieve this, the project involves several key steps:

1. **Data Collection and Preprocessing:** To make sure historical weather data is suitable for modelling, it is gathered and pre-processed. These covers normalizing numerical features, encoding categorical data into numerical formats, and using imputation to handle missing values. The dataset is cleansed to get rid of features and columns that aren't necessary for predicting.
2. **Model Development:** Two prediction models are created, one with Random Forest and the other with LSTM. The sequential weather data is sent into the LSTM model, which is

meant to extract temporal dependencies and identify patterns that might point to future weather. Multiple decision trees make up the Random Forest model, which is trained to comprehend intricate correlations between many meteorological data and enhance forecast performance.

3. **Model Evaluation:** Many criteria are used to assess the performance of both models, including F1-score, accuracy, precision, and recall. This assessment aids in determining how well each algorithm predicts rainfall and offers a framework for contrasting the two methods. To make sure the models are reliable and applicable to a wide range of situations, cross-validation methods are used.

4. **Results and Analysis:** The findings show that, in comparison to conventional techniques, both Random Forest and LSTM models greatly improve predicting accuracy. When it comes to forecasting changes in the weather over time, the LSTM model shows a great capacity to identify temporal patterns and trends in the data. Because of its ensemble approach, the Random Forest model offers great accuracy and manages a wide range of features well. The comparison analysis provides insights into the advantages of each algorithm for weather forecasting while highlighting the drawbacks and strengths of each method.

5. **Conclusion and Recommendations:** The study comes to the conclusion that incorporating machine learning algorithms like Random Forest and LSTM can significantly increase the accuracy of weather forecasting. The results imply that even greater outcomes might be achieved by combining these strategies or by utilizing their complimentary qualities. To improve prediction even more, future research may investigate hybrid models, include other data sources, or apply these methods to various forecasting jobs.

In summary, this effort advances the area of weather forecasting by illuminating the ways in which sophisticated machine learning methods might be used to the problems associated with weather prediction. The application of Random Forest and LSTM algorithms offers a fresh method for improving forecasting accuracy and gives insightful information for future studies and real-world meteorological applications.

Table of Contents

CHAPTER NO.	CHAPTER NAME	PAGE NO.
1	Introduction	
	1.1 Background Information on Weather Forecasting	
	1.2 Importance and Applications of Accurate Weather Predictions	
	1.3 Overview of Machine Learning, Random Forest, and LSTM in Weather Forecasting	
2	Literature Review	
	2.1 Overview of Existing Methods for Weather Forecasting	
	2.2 Traditional Weather Forecasting Methods	
	2.3 Random Forest in Weather Forecasting	
	2.4 Long Short-Term Memory (LSTM) Networks in Weather Forecasting	

	2.5 Comparison Between Traditional Methods and Machine Learning Approaches	
	2.6 Key Studies and Advancements in the Field	
3	Problem Statement	
	3.1 Defining the Problem	
	3.2 Challenges in Accurate Weather Forecasting	
	3.3 Approach to Solving the Problem	
4	Methodology	
	4.1 Data Collection	
	4.2 Data Preprocessing	
	4.3 Model Architecture	
	4.4 Training and Testing	

5	Implementation	
	5.1 Weather forecasting using Random Forest	
	5.2 Weather forecasting using LSTM	
6	Result And Discussion	
	6.1 Overview of the model	
	6.2 Random Forest Model	
	6.3 LSTM Model	
	6.4 Comparative Analysis	
7	Conclusion and future work	
	7.1 Summary of the project	
	7.2 Performance Evaluation and Insights	
	7.3 Comparative Analysis	
	7.4 Practical Implication	
	7.5 Limitations and future works	
	7.6 Concluding Remarks	
8	References	

CHAPTER 1

1.INTRODUCTION

1.1 Background Information on Weather Forecasting

Historical Development of Weather Forecasting

Systematic weather observation started in the 17th century with the creation of weather instruments. Evangelista Torricelli's creation of the barometer in 1643 made it possible to monitor atmospheric pressure, which was essential for forecasting changes in the weather. (*E. Torricelli, Opera Geometrica, 1644*). Similar to this, Galileo Galilei and associates invented the thermometer in the early 17th century to monitor temperature—another crucial factor in weather prediction.

The development of numerical weather prediction (NWP) models in the 20th century was the true breakthrough in weather forecasting. These models, which were first presented by Lewis Fry Richardson in 1922, simulated atmospheric phenomena using mathematical equations derived from physical laws. (*L. F. Richardson (1922)*). Weather Forecasting using Numerical Methods but Richardson's approach was constrained by the available computing power at the time. NWP models were impractical until the 1950s, when digital computers were developed, making it possible for meteorologists to process enormous volumes of data and generate forecasts that were more precise.

Radar systems, weather stations, and satellite data have all been included into more complex NWP models throughout the years. Advanced global models built by organizations such as the European Centre for Medium-Range Weather Forecasts (ECMWF) and the National Oceanic and Atmospheric Administration (NOAA) can predict weather with amazing accuracy. These models still have limitations, though, especially when it comes to coping with the chaotic structure of the environment and forecasting localized weather events.

Challenges in Traditional Weather Forecasting

Traditional weather forecasting techniques have many drawbacks despite their improvements. One of the biggest obstacles is the atmosphere's chaotic quality, which is sometimes called the

"butterfly effect." Meteorologist Edward Lorenz first proposed this idea in the 1960s. It implies that slight variations in starting conditions can produce wildly disparate results in weather forecasts. (**E. N. Lorenz, 1963**). Nonperiodic deterministic flow.) Because of its intrinsic unpredictability, long-term forecasts are less accurate and it is more challenging to forecast abrupt weather changes, such the development of a storm.

The computational complexity of NWP models is another difficulty. In order to solve the intricate mathematical equations that explain atmospheric dynamics, these models demand a significant amount of processing power. Because of this, producing precise projections requires a lot of time and computing resources, especially when doing it globally. Forecasts are frequently delayed as a result, which can be especially troublesome for short-term projections where timely information is essential.

Furthermore, it is frequently difficult for conventional models to forecast extreme weather phenomena like hurricanes, tornadoes, and flash floods. Localized phenomena that are challenging to represent using the wide geographical and temporal scales employed in global NWP models have an impact on these events. Hurricane Katrina in 2005 and Hurricane Harvey in 2017 are two examples of how models frequently fail to foresee the rapid strengthening of hurricanes.

A rising number of people are interested in combining cutting-edge computational approaches, especially machine learning, with conventional forecasting methods in light of these difficulties. Machine learning has the ability to improve weather forecasting's efficiency and accuracy by utilizing big datasets and complex algorithms.

1.2 Importance and Applications of Accurate Weather Predictions

Agriculture

Since the agricultural industry depends heavily on weather, accurate weather forecasting is essential. For the purpose of planting, watering, controlling pests, and harvesting, farmers rely on weather forecasts. To guarantee that crops receive enough water during their crucial growth stages, farmers might, for instance, organize their planting schedules more effectively by anticipating the

start of the rainy season. In a similar vein, farmers can reduce damage by covering crops or utilizing frost protection devices in response to frost warnings.

Significant losses in agriculture can result from inaccurate weather forecasts. Unexpected droughts or an abundance of rainfall, for example, might destroy crops, resulting in lower yields and financial losses for farmers. According to a study done by the Food and Agriculture Organization (FAO), unfavourable weather causes 26% of crop losses worldwide (*Food and Agriculture Organization (FAO), 2018*).

Weather forecasting is essential for guaranteeing food security even beyond its immediate economic effects, especially in areas that are vulnerable to catastrophic weather occurrences. For instance, precise rainfall forecasts are crucial to sustaining agricultural productivity and averting food shortages in sub-Saharan Africa, where agriculture is predominantly rain-fed (*Food and Agriculture Organization (FAO), 2018*).

Weather forecasting is essential for guaranteeing food security even beyond its immediate economic effects, especially in areas that are vulnerable to catastrophic weather occurrences. For example, precise rainfall forecasts are crucial to sustaining agricultural productivity and averting food shortages in sub-Saharan Africa, where agriculture is predominantly rain-fed.

Transportation delays brought on by bad weather can have serious economic repercussions. According to FAA estimates, weather-related delays account for 70% of all air traffic delays in the US, resulting in billions of dollars of annual losses for the airline sector (**Federal Aviation Administration (FAA), 2021**). Unexpected storms can result in the loss of cargo and even ships in the maritime industry, as demonstrated by Hurricane Joaquin's 2015 sinking of the El Faro.

Accurate weather forecasting is also important for road safety. Unfavourable weather, such as snow, ice, and a lot of rain, can make driving dangerous and raise the chance of accidents. For instance, according to data from the National Highway Traffic Safety Administration (NHTSA), weather-related incidents cause about 6,000 fatalities and over 445,000 injuries annually in the United States, or 21% of all vehicle accidents. 2021: *National Highway Traffic Safety Administration (NHTSA)* documents.

Disaster Management

In disaster management, weather forecasting is an essential tool, especially when it comes to anticipating and reducing the effects of natural catastrophes like hurricanes, floods, and wildfires. Authorities can limit the loss of life and property by planning evacuations, issuing early warnings, and mobilizing resources in response to timely and accurate meteorological predictions.

For example, forecasting Hurricane Sandy's impact in 2012 enabled the evacuation of approximately 375,000 people in New York City, considerably decreasing the possible death toll. (*National Oceanic and Atmospheric Administration, 2012*) In similar fashion, early warning systems and flood forecasts have played a critical role in averting significant fatalities in areas vulnerable to flash floods, such as Bangladesh and the Mississippi River basin in the United States.

However, the effectiveness of disaster management programs is directly impacted by how accurate weather forecasts are. As seen by Hurricane Maria in 2017, when the storm's intensity was overstated, causing a delay in emergency response and significant damage in Puerto Rico, poor forecasts can lead to inadequate planning (*National Hurricane Centre, 2018*).

Energy Sector

Accurate weather forecasts are critical to the energy sector, especially given the growing reliance on renewable energy sources. Since weather patterns affect the production of wind, solar, and hydropower, precise forecasting is crucial to maintaining a balance between supply and demand.

For example, wind direction and speed have a direct impact on wind power generation. Grid operators can predict power output and modify supply to fit demand with the help of accurate wind forecasts. In areas where wind energy is widely used, like certain portions of Europe and the United States, this is essential to preserving grid stability and averting blackouts. (*Wang, L., & Hu, Z., 2020*).

Similar to this, solar radiation levels and cloud cover have an impact on solar power generation. Energy providers can anticipate variations in power output and make the required adjustments with the help of accurate solar forecasts. For instance, precise solar forecasting helps avoid energy shortages during periods of peak demand in California, where solar power constitutes a sizable component of the energy mix. 2020; (*California ISO, 2020*).

Precise predictions of precipitation and snowmelt are critical to the hydropower industry's ability to manage water resources and guarantee reliable power production. For example, precise forecasts aid in reservoir level management and power generating optimization in the Pacific Northwest region of the United States, where hydropower accounts for a sizable amount of the electricity generated. (*California ISO, 2020*).

Daily Life

Weather forecasts have a significant impact on daily decision-making, influencing everything from outfit choices to schedules of activities. For instance, judgments on whether to bring an umbrella, wear a coat, or postpone outdoor activities are influenced by precise forecasts of rain, snow, or extremely high temperatures. Despite their seeming insignificance, these choices have important social and economic ramifications.

For example, the weather can have an impact on store sales and consumer behaviour. According to a British Retail Consortium survey, sales of winter apparel fell in 2015 due to unexpectedly warm weather, yet sales of coats and heating products increased in 2013 due to extremely cold weather. (*British Retail Consortium, 2016*).

Emergency services and public safety are also impacted by weather. Accurate forecasting, for instance, might lead to public health recommendations during heatwaves to prevent heat-related illnesses, especially in vulnerable groups like the elderly and people with pre-existing medical disorders. (*World Health Organization, 2018*). In a similar vein, weather forecasts can help with

planning during extreme cold spells so that hypothermia and other cold-related health problems are avoided.

1.3 Overview of Machine Learning, Random Forest, and LSTM in Weather Forecasting

[Introduction to Machine Learning in Weather Forecasting](#)

The goal of machine learning (ML), a branch of artificial intelligence (AI), is to create algorithms that can learn from and forecast data. In contrast to traditional programming, which follows predetermined instructions, machine learning algorithms employ statistical methods to find patterns in data and enhance their functionality over time. Because of its capacity for learning and adaptation, machine learning (ML) is especially well-suited for tasks like time-series forecasting, pattern identification, and anomaly detection—all of which are critical to weather prediction.

Neural networks were employed in the field of climate modelling, one of the first uses of machine learning (ML) in meteorology, to approximate the results of intricate models. Since then, ML has been applied to a wide range of forecasting jobs, such as identifying severe weather phenomena like hurricanes and tornadoes, as well as predicting temperature and precipitation. When combined with standard NWP models, ML models can improve their accuracy by removing biases and/or supplying localized forecasts in areas where traditional models are not as accurate. (*Hochreiter, S., & Schmidhuber, J., 1997*).

Random Forest in Weather Forecasting

Since Random Forest (RF) can handle huge datasets and is resilient against overfitting, it is a common ensemble learning technique used in weather forecasting. Leo Breiman created RF in 2001. It works by building several decision trees during training and producing the mean prediction (for regression) or the mode of the classes (for classification) of each individual tree. (*Breiman, L., 2001*).

When evaluating big meteorological datasets for weather forecasting purposes, RF is very helpful because it can be used to predict a variety of weather factors, including wind speed, temperature, and humidity. For example, RF can be used to forecast the chance of precipitation by examining

past meteorological information, which includes elements like pressure, temperature, and dew point.

The ability of RF to handle nonlinear interactions between variables, which are typical in meteorological data, is one of its main advantages in weather forecasting. Furthermore, RF can offer feature importance estimations, assisting meteorologists in determining which factors have the greatest bearing on forecasting particular weather events.

Research has demonstrated through case studies that RF can greatly increase predicting accuracy. For instance, RF performed better than conventional statistical models in research on the forecasting of heavy rainfall in South Korea, yielding forecasts that were more accurate and dependable (*Lee, J. et al., 2019*). Comparably, temperature anomaly prediction—which is essential for spotting heatwaves and cold snaps—has been enhanced by RF use (*Lee, J. et al., 2019*).

Long Short-Term Memory (LSTM) in Weather Forecasting

Recurrent neural networks (RNNs) of the Long Short-Term Memory (LSTM) type function particularly well for time-series forecasting, which makes them perfect for weather prediction. LSTM networks, which were first presented by Sepp Hochreiter and Jürgen Schmidhuber in 1997, are intended to capture long-term dependencies in sequential data and solve the vanishing gradient issue that typical RNNs have. (*Hochreiter, S., & Schmidhuber, J., 1997*).

LSTM networks are particularly good at capturing temporal dependencies in weather data, like the diurnal cycle of air pressure or seasonal patterns in temperature, when it comes to weather forecasting. Additionally, LSTM networks perform well while processing inadequate and noisy data, which is frequently the case with meteorological observations.

Predicting temperature trends and seasonal fluctuations is one of the main uses of LSTM in weather forecasting. One use of LSTM networks is the analysis of historical temperature data, together with other pertinent variables such as humidity and wind speed, to forecast daily

temperature changes. Planning in the fields of public health, energy management, and agriculture will benefit greatly from these forecasts.

Analyses of specific cases have shown how well LSTM works to increase weather forecast accuracy. For example, LSTM performed better than conventional autoregressive models in a study on temperature forecasting in Beijing, yielding more accurate short- and long-term predictions (*Zhang, Y. et al., 2020*). In a similar vein, wind speed—a crucial factor in the generation of wind energy—has been highly accurately predicted by LSTM networks (*Wang, L., & Hu, Z., 2020*).

Combining Random Forest and LSTM for Enhanced Forecasting

Because of their complementing advantages, Random Forest and LSTM networks perform well together to increase the accuracy of weather forecasts. LSTM networks are skilled at capturing temporal dependencies in sequential data, but RF is best at managing big datasets and recognizing significant features. These two techniques can be used to build hybrid models that take advantage of each other's advantages.

To prepare meteorological data for time-series forecasting, a hybrid model could, for instance, employ RF to identify important variables and reduce noise before feeding the data into an LSTM network. More precise forecasts may result from this method, especially for intricate weather occurrences including both temporal and spatial patterns. (*Wang, L., & Hu, Z., 2020*).

The potential of these hybrid models has been demonstrated by recent studies. In one work on wind speed prediction, for example, RF and LSTM were combined to provide forecasts that were more accurate than those produced by either technique alone (*Wang, L., & Hu, Z., 2020*). The adaptability and efficacy of this strategy have been demonstrated by the application of hybrid models to forecast temperature, rainfall, and other meteorological variables (*Lee, J. et al., 2019*).

1.4 Project Objective

1. **Develop a machine learning-based weather forecasting system** utilizing Random Forest (RF) and Long Short-Term Memory (LSTM) algorithms.
2. **Leverage historical weather data** to build predictive models that can accurately forecast future weather conditions.
3. **Model Comparison:** Evaluate the performance of Random Forest and LSTM models in predicting weather patterns, highlighting their strengths and weaknesses in forecasting accuracy.

CHAPTER 2

2. Literature Review

2.1 Overview of Existing Methods for Weather Forecasting

Over the ages, weather forecasting has changed dramatically, moving from crude observations to extremely complex models that make use of cutting-edge computing tools. The first forecasting techniques relied on empirical laws and trends that were noticed over time. For instance, cloud formations, wind directions, and other natural indications were used by farmers and sailors to forecast changes in the weather. Although helpful, these conventional methods have limitations in terms of precision and reach.

A significant breakthrough in weather forecasting occurred in the 20th century with the development of statistical methodologies. Historical weather data was analysed to find trends that can be utilized to forecast future conditions using statistical models like autoregressive and linear regression. These models laid the foundation for more sophisticated methods by being among the first to apply mathematical concepts to weather forecasting.

In the middle of the twentieth century, a novel method called numerical weather prediction (NWP) was developed. NWP models simulate the atmosphere using mathematical equations that describe physical processes like fluid dynamics and thermodynamics. They were developed by organizations like the European Centre for Medium-Range Weather Forecasts (ECMWF) and were first developed by scientists like Lewis Fry Richardson (*Richardson, L. F., 1922*). In order to forecast the future states of the atmosphere, these models divide the atmosphere into a grid and solve the equations at each grid point.

NWP models are successful, although they have a few drawbacks. Due to the "butterfly effect," which is the chaotic character of the environment, even tiny beginning condition mistakes can cause large long-term variations in forecast accuracy (*Lorenz, E. N., 1963*). Furthermore, processing the enormous volumes of data required in NWP models requires supercomputers because to their high computational needs. Because of this, NWP models are quite useful for short- to medium-range forecasts, but as time goes on, their accuracy decreases.

Hybrid strategies that integrate machine learning (ML) and conventional NWP methods have been created in response to these difficulties. In situations when standard methods might not be as successful, these approaches seek to increase the accuracy of localized forecasts by addressing biases in NWP models. Large datasets may be processed quickly by ML models, which can also spot patterns that physical modelling might miss. This has created new opportunities to improve weather forecast accuracy and dependability.

2.2 Traditional Weather Forecasting Methods

Statistical and dynamical approaches can be used to broadly categorize traditional weather forecasting techniques. The assumption of statistical approaches is that future weather patterns will match past ones, based on historical data. These approaches include the autoregressive integrated moving average (ARIMA) method (Box, *G.E.P.*, *Jenkins, G.M.*, *Reinsel, G.C.*, & *Ljung, G.M.*, *2015*), which predicts the links between past and future values of a time series. Applications ranging from forecasting seasonal temperature fluctuations to estimating rainfall totals have made extensive use of ARIMA and related models.

The simplicity and convenience of use of statistical models is one of their main benefits. In comparison to dynamical models, they need less processing power and may produce forecasts for specific weather parameters that are reasonably accurate in a timely manner. Nevertheless, they find it difficult to anticipate major weather events or to record abrupt changes in atmospheric conditions since they rely so heavily on historical data. Furthermore, these models use the assumption that the fundamental statistical characteristics of the meteorological data won't change over time, which isn't necessarily the case because of other variables like climate change.

An alternative approach is taken by dynamic methods, symbolized by NWP, which solve the fluid dynamics equations governing the behaviour of the atmosphere. The foundation of NWP models, such as the Global Forecast System (GFS) and the European ECMWF model, is the conservation of mass, energy, and momentum—three fundamental laws of physics. By modeling the interactions between different atmospheric factors, such as temperature, pressure, humidity, and wind speed, these models are able to produce precise forecasts.

Although NWP models are extremely complex, their setup requires enormous quantities of data and is computationally demanding. The quality of the original data and the grid resolution utilized in the simulations have a significant impact on how accurate these models are. Finer details of atmospheric processes can be captured by high-resolution models, but they require more processing power and time to run. One of the most important factors in operational forecasting is the trade-off between computational viability and accuracy.

The shortcomings of deterministic NWP models have been addressed in recent years with the introduction of ensemble forecasting. Several simulations with marginally different initial circumstances are run as part of ensemble forecasting in order to account for the uncertainty in the state of the atmosphere. Forecasters are able to determine probable extreme weather events and gauge forecast confidence by examining the distribution of the ensemble members. In contemporary meteorology, ensemble forecasting has become standard procedure because it improves decision-making and offers insightful information about prediction uncertainty.

2.3 Random Forest in Weather Forecasting

The ensemble learning technique known as Random Forest (RF) has become well-known in weather forecasting because of its resilience, adaptability, and capacity to manage huge, intricate datasets. In order to generate predictions, RF builds several decision trees during training and aggregates the results. By using an ensemble technique, the likelihood of overfitting is decreased and the model's capacity for generalization is enhanced. (*Breiman, L., 2001*).

When it comes to weather forecasting, RF works especially well for jobs like wind speed, precipitation, and temperature prediction that require a lot of historical weather data. The model is appropriate for a variety of meteorological applications since it can handle a large range of input variables, including both numerical and categorical data.

The capacity of RF to identify nonlinear correlations between variables—which are typical in meteorological data—is one of its main advantages. For instance, the correlation between humidity and temperature is frequently nonlinear, with certain combinations producing more intense weather. Forecasts produced by RF are more accurate because it is a more effective modeler of these intricate relationships than linear models.

Numerous studies have shown that RF is useful for a variety of weather forecasting applications. For example, RF was used to evaluate historical meteorological data, such as temperature, humidity, and atmospheric pressure, to forecast rainfall levels in a study on heavy rainfall prediction in South Korea. According to the study, RF performed better than conventional statistical models and produced forecasts that were more accurate and trustworthy, particularly for instances of intense rainfall. (*Lee, J. et al., 2019*).

Temperature anomalies, or departures from the expected temperature range, were predicted using RF in a different study. Temperature anomalies are essential for predicting meteorological phenomena like heat waves and cold snaps, which can have major effects on energy use, agriculture, and public health. In comparison to conventional approaches, the study demonstrated that RF could capture the intricate interactions between many atmospheric factors, resulting in more accurate predictions of temperature anomalies. (*Lee, J. et al., 2019*).

Additionally, RF offers estimates of feature importance, which are helpful in determining which factors have the greatest bearing on the forecast of particular weather conditions. With the help of this tool, meteorologists can improve their forecasting models by learning more about the elements influencing specific weather patterns.

2.4 Long Short-Term Memory (LSTM) Networks in Weather Forecasting

Time-series forecasting is a particularly good fit for Long Short-Term Memory (LSTM) networks, a specific kind of recurrent neural network (RNN) intended to simulate sequential data. Traditional RNNs struggle with the vanishing gradient issue, which LSTM networks solve, making it possible for them to identify long-term dependencies in data. (*Hochreiter, S., & Schmidhuber, J., 1997*).

LSTM networks have been used in the field of weather forecasting for a variety of time-series prediction tasks, such as wind speed estimation, precipitation prediction, and temperature forecasting. The capacity of LSTM networks to learn from sequential data and recognize patterns and trends over time—which are essential for precise weather forecasts—is one of their main advantages.

LSTM networks, for instance, have been used to forecast daily temperature changes by examining previous temperature data in conjunction with other pertinent factors like wind speed and humidity. Numerous fields, including public health, agriculture, energy management, and disaster planning, can benefit from these forecasts. LSTM networks fared better than conventional autoregressive models in research on temperature forecasting in Beijing, yielding more accurate short- and long-term predictions. (*Zhang, Y. et al., 2020*).

Additionally, LSTM networks perform well while processing inadequate and noisy data, which is frequently the case with meteorological observations. Many things, including inconsistent measurements, missing numbers, and sensor faults, can have an impact on weather data. Because LSTM networks are resistant to these problems, they can generate accurate forecasts even in cases when the input data is incomplete.

Wind speed prediction is a crucial use of LSTM networks in weather forecasting and is essential for the production of renewable energy, especially in wind farms. Precise predictions of wind speed lessen the hazards associated with erratic wind patterns and improve the planning and functioning of wind energy installations. LSTM networks were able to capture the temporal dependencies in wind speed data in wind speed prediction research, producing forecasts that were more accurate than those made using conventional approaches. (*Wang, L., & Hu, Z., 2020*).

To improve predicting accuracy, LSTM networks have also been employed in conjunction with other ML models, such as RF. With LSTM concentrating on time-series prediction and RF handling feature selection and data preparation, these hybrid models combine the best features of both methodologies. It has been demonstrated that combining these techniques improves forecast performance in a number of applications, including wind speed estimation and the prediction of temperature and precipitation.

2.5 Comparison Between Traditional Methods and Machine Learning Approaches

The area has made great strides over the past few decades, as evidenced by the comparison between modern machine learning algorithms and classic weather predicting methods. Conventional techniques, including statistical models and NWP, base their predictions on known physical laws and past data. These approaches have limits, especially when it comes to capturing the intricate,

nonlinear interactions between atmospheric factors and predicting catastrophic weather occurrences, even though they have been successful in generating general forecasts.

Compared to conventional techniques, machine learning techniques like RF and LSTM have a number of advantages. Large, high-dimensional datasets can be handled by RF models, which can also represent intricate, nonlinear interactions between variables. This makes them especially useful for forecasting unique weather occurrences, like intense downpours or unusually high or low temperatures, which can be difficult for conventional models to properly represent.

On the other hand, temporal dependencies in time-series data are well-represented by LSTM networks, which are excellent at modelling sequential data. Since many weather patterns, like seasonal cycles and daily fluctuations, exhibit substantial temporal correlations, this capacity is essential for good weather prediction. Compared to conventional methods, LSTM networks can produce more accurate and dependable forecasts by learning these patterns from past data, especially for short- and medium-term predictions.

But there are drawbacks to machine learning models as well. They can be computationally demanding and require a lot of data for training. Furthermore, compared to conventional techniques, these models are harder to comprehend due to their "black box" nature. When decision-making depends on knowing the underlying mechanisms of weather patterns, this lack of interpretability may be problematic.

The future of weather forecasting looks quite promising when machine learning is combined with conventional techniques, despite these obstacles. Researchers can create forecasting models that are more accurate and dependable while addressing the shortcomings of each technique by integrating the advantages of both methodologies. With the increasing growth of processing power and availability of additional meteorological data, this hybrid technique is expected to gain greater significance.

2.6 Key Studies and Advancements in the Field

Weather forecasting has advanced recently, with an emphasis on creating hybrid models that combine the best features of machine learning and conventional techniques. These models leverage

the complementing capabilities of various methodologies in an attempt to improve forecast accuracy.

One such development is the integration of RF and LSTM to increase wind speed prediction accuracy. A hybrid RF-LSTM model was created in a Wang and Hu (2020) study to forecast wind speed in a wind farm. While the LSTM component of the model concentrated on simulating the temporal relationships in the wind speed data, the RF component was utilized to identify the most pertinent features from the input data. In comparison to standalone RF and LSTM models, the hybrid model performed better and produced more accurate and consistent wind speed forecasts. (*Wang, L., & Hu, Z., 2020*).

The combination of machine learning and ensemble forecasting methods is another significant development. In classical weather forecasting, ensemble forecasting—which entails conducting many simulations with varied initial conditions—has been a regular procedure. The accuracy and dependability of the forecasts can be increased by researchers by including machine learning models into the ensemble framework. For instance, biases in the ensemble members can be fixed by machine learning models, improving the accuracy of extreme weather event forecasts.

Another developing field of study in weather forecasting is the use of deep learning methods like generative adversarial networks (GANs) and convolutional neural networks (CNNs). Even though RF and LSTM are the main topics of this analysis, it's important to remember that other cutting-edge methods are being investigated for applications like precipitation nowcasting, where the objective is to forecast short-term precipitation patterns with high spatial and temporal precision. These methods offer more precise and in-depth predictions in real-time, which has the potential to completely transform weather forecasting.

The combination of conventional techniques with machine learning and deep learning approaches is expected to be vital in improving our capacity to anticipate and address weather-related issues as the field of weather forecasting develops. The next generation of weather forecasting systems will provide more precise and dependable predictions for a variety of applications. This will be made possible by the continuous development of hybrid models, the use of ensemble approaches, and the investigation of novel machine learning algorithms.

CHAPTER 3

1. Problem Statement

A major issue affecting many industries, including agriculture, emergency preparedness, transportation, and everyday human activity, is accurate weather forecasting. Because weather systems are highly dynamic and non-linear, predicting weather conditions, especially rainfall, remains a complex and challenging process despite substantial breakthroughs in meteorological research and technology. This study uses machine learning techniques to handle the problem of rainfall prediction, which is a crucial aspect of weather forecasting.

3.1 Defining the Problem

This project's main goal is to solve the challenge of reliably forecasting rainfall from past weather data. Predicting rainfall is crucial for many uses, including managing crops, preventing flooding, and allocating resources. The intricacy of weather patterns can be difficult for statistical models and physical simulations, which are frequently used in traditional weather forecasting techniques. These approaches might not adequately represent the complex interrelationships between many meteorological factors or be able to adjust to a changing climate.

To tackle this problem, the project explores the use of two advanced machine learning algorithms:

1. **Random Forest:** A potent ensemble learning technique that builds several decision trees and combines them to reduce overfitting and increase prediction accuracy. When it comes to managing structured data and capturing intricate relationships between input features, Random Forest excels.
2. **Long Short-Term Memory (LSTM) Networks:** A kind of recurrent neural network (RNN) with a long retention period that is especially made to handle sequential data. For applications involving time series data, like weather forecasting, where previous data points might have a big impact on future predictions, LSTMs are a good fit.

3.2 Challenges in Accurate Weather Forecasting

Accurate weather forecasting, especially rainfall prediction, is fraught with challenges that make it a difficult problem to solve:

1. **Complexity of Weather Systems:** Numerous elements, such as temperature, humidity, wind patterns, and atmospheric pressure, have an impact on weather systems. It might be challenging to model these components' complicated, frequently non-linear interactions using conventional methods. With its ensemble of decision trees, the Random Forest algorithm models these interactions, which helps alleviate this complexity.
2. **Temporal Dependencies:** Because weather data is time-dependent and sequential, previous weather conditions may have an impact on current and future situations. These temporal relationships are frequently difficult for traditional machine learning methods to accurately capture. Because LSTM networks are built to handle sequential data, they are very helpful in weather forecasting, where it is essential to comprehend the temporal correlations between variables.
3. **Data Quality and Availability:** Measurement mistakes, missing data points, and different data gathering techniques are just a few of the variables that can cause weather data to be noisy, inaccurate, and inconsistent. Data cleaning, imputation, and feature engineering are a few of the strong data preparation approaches that must be used to ensure data quality, as it is a major difficulty.
4. **Model Generalization:** Developing a model that accurately extrapolates to unknown data is crucial to accurate forecasting. One common issue is overfitting, which occurs when a model performs well on training data but badly on fresh data. This project will make use of the complimentary qualities of Random Forest, which has built-in techniques to avoid overfitting, and LSTM, which can capture long-term dependencies.
5. **Computational Resources:** Large amounts of computational time and resources are needed to train intricate models like LSTMs. An essential factor in the project is striking a balance between the complexity of the model and the resources at hand.
6. **Feature Selection and Engineering:** Building successful predictive models requires first sifting through the massive amount of meteorological data to find and pick the most

pertinent elements. By generating new features that effectively capture major patterns in the data, proper feature engineering can dramatically improve the performance of both Random Forest and LSTM models.

3.3 Approach to Solving the Problem

This project proposes a systematic approach to address these challenges and improve the accuracy of rainfall prediction:

1. **Data Preprocessing:** To guarantee that the dataset is of the highest quality and prepared for modelling, the first stage is cleaning the data, addressing missing values, and encoding categorical variables.
2. **Model Development:** Two models will be developed as part of the project: an LSTM model and a Random Forest model. In order to identify the patterns connected to rainfall occurrences, each model will be trained using historical meteorological data.
3. **Model Evaluation and Comparison:** We will compare the two models' performance using metrics like F1-score, recall, accuracy, and precision. The project compares different models to see which algorithm performs better in terms of prediction when it comes to rainfall forecasting.
4. **Optimization and Fine-Tuning:** To maximize performance, hyperparameter adjustment will be applied to both models. To make sure the models perform well when applied to new data, cross-validation and other techniques will be employed.
5. **Feature Engineering:** To increase the models' ability to forecast the future, the research will investigate a variety of feature engineering strategies. This could entail adding new data sources or developing new features from already-existing data. By using this strategy, the project hopes to create reliable and accurate models that may greatly enhance rainfall predictions, which would ultimately lead to more efficient decision-making across a range of industries.

CHAPTER 4

4. Methodology

4.1 Data Collection

The relevance and quality of the data used for training and testing have a significant impact on the accuracy of machine learning models. Historical weather data was gathered for this research, with an emphasis on important meteorological factors like temperature, humidity, wind speed, and precipitation. Reputable weather stations and internet meteorological databases provided the multi-year dataset. (Torricelli, 1644).

The present study employed a CSV dataset, which is a widely used and adaptable file format for data processing and archiving (Han et al., 2011). There were several columns in the data, each of which represented a distinct weather aspect. Important characteristics comprised:

- **Temperature:** With figures expressed in degrees Celsius, the daily minimum, maximum, and average are provided.
- **Humidity:** The percentage that represents the amount of moisture in the air.
- **Wind Speed:** Information on the wind speed at different times, expressed in kilometres per hour.
- **Precipitation:** Millimetres are used to measure the amount of rainfall that falls during a specific time frame.
- **Other Variables:** There were also more elements including visibility, barometric pressure, and wind direction.

This dataset provided the required inputs for predictive modelling and formed the basis for training the Random Forest and LSTM models. (Cheng et al., 2017).

4.2 Data Preprocessing

In order to guarantee that the dataset is clear, consistent, and prepared for model training, data preparation is an essential step. The preprocessing actions listed below were carried out:

1. **Data Cleaning:** Some unnecessary columns from the dataset, including "_precipm" and "datetime_utc," were removed because they weren't useful for the forecasting purpose. By

eliminating these columns, the dataset was made simpler and the study was concentrated on pertinent aspects. (Pedregosa et al., 2011).

2. **Handling Missing Values:** In large datasets, missing data is a prevalent problem that can seriously affect how well machine learning model's function. In this project, the sklearn library's SimpleImputer was used to fill in the missing values in numerical columns. To ensure that the dataset was preserved and that all features were used during the model training process, the imputer substituted the mean of the corresponding column for any missing values. (Wooldridge, 2013).
3. **Categorical Variable Encoding:** Categorical variables, such as weather conditions (e.g., "Clear," "Cloudy," "Rainy") were included in the dataset. **LabelEncoder** was utilized to transform this categorical information into a numerical representation by allocating a distinct integer to every category. This step was essential to ensuring that the data could be processed and interpreted by the machine learning algorithms. (Friedman et al., 2001).
4. **Feature Selection and Normalization:** Feature selection is crucial for improving model performance since it gets rid of redundant or unnecessary data. Only the most important traits remained after preprocessing. The next step was to normalize the numerical features such that their values fell into a predetermined range, usually between 0 and 1. For the LSTM model, which is sensitive to the size of numerical inputs, normalization was very vital. (LeCun et al., 2012).
5. **Data Splitting:** Training and testing sets were created using the preprocessed dataset in order to assess the model's performance. 30% of the data was usually set aside for testing, and the remaining 70% was used for training. This division guarantees that the model is assessed using unobserved data, offering an accurate assessment of its forecasting skills. (Kohavi, 1995).

4.3 Model Architecture

In this study, two different machine learning models were used: **Random Forest** classifiers and **Long Short-Term Memory (LSTM)** networks. These models were selected due to their ability to handle complicated feature interactions and sequential data in a complementing way.

1. LSTM Model Architecture:

- **Input Layer:** The pre-processed sequential weather data, with each time step representing a specific day's weather observations, is sent into the input layer of the LSTM model. (**Hochreiter & Schmidhuber, 1997**).
- **LSTM Layers:** The network that processes the time series data by preserving and updating internal states (memory cells) that record long-term dependencies is known as the LSTM layers, and it is this heart of the model. The number of units (neurons) in each of the architecture's numerous LSTM layers impacts the model's ability to recognize temporal patterns. (**Graves, 2012**).
- **Dropout Layers:** Dropout layers were inserted in between LSTM layers to stop overfitting. During training, these layers arbitrarily remove a portion of the connections between neurons, which forces the model to improve its generalization. (**Srivastava et al., 2014**).
- **Dense Layers:** Dense (completely linked) layers were added after the LSTM layers in order to aggregate the information retrieved by the LSTM units and generate the final output. The learned characteristics are mapped to the output space via the thick layers. (**Goodfellow et al., 2016**).
- **Output Layer:** The final weather forecasts, including the likelihood of rain or other weather events, are produced by the output layer.

2. Random Forest Model Architecture:

- **Bootstrap Sampling:** Using several dataset subsets, the Random Forest ensemble approach creates several decision trees. These subsets are produced by bootstrap sampling, which involves randomly selecting replacement samples from the dataset. Because each tree is trained on a different subset, the model is more robust overall. (**Breiman, 2001**).
- **Decision Trees:** Recursively dividing the dataset according to the most informative features results in the construction of each decision tree in the Random Forest. The splits are selected to optimize a certain criterion, like Gini impurity or information gain. This procedure keeps going until every tree reaches maturity. (**Quinlan, 1986**).

- **Ensemble Voting:** To arrive at the ultimate prediction, the Random Forest model combines the forecasts from each decision tree. Majority voting is usually used for classification tasks, like rainfall prediction, where the class with the most votes from the trees is selected as the final result. (Dietterich, 2000).

4.4 Training and Testing

For the purpose of creating reliable and accurate models, the training and testing stages were essential. After each model was trained on the training set, it was assessed using a variety of performance measures on the testing set.

1. LSTM Training:

- **Training Process:** Backpropagation through time (BPTT) was used to train the LSTM model. This entails utilizing an optimizer such as Adam to update the model parameters by computing the gradient of the loss function with respect to those parameters. Several epochs were used for the training, with each epoch denoting a full run through the training dataset. (Rumelhart et al., 1986).
- **Early Stopping:** We used early stopping to avoid overfitting. This method prevents overtraining, which could result in overfitting, by tracking the model's performance on a validation set during training and stopping training when the performance stops improving. (Prechelt, 1998).

2. Random Forest Training:

- **Training Process:** In order to train the Random Forest model, many decision trees were built, each using a distinct bootstrap sample of the data. Growing each tree until it was fully developed—that is, when every leaf node held data from a single class or a predetermined minimum number of samples—was the training process. (Breiman, 2001).
- **Out-of-Bag Error:** The model's performance during training was assessed using an out-of-bag (OOB) error estimate. This estimate, which offers an objective

assessment metric, is obtained from the subset of data for each tree that was excluded from the bootstrap sample. (Breiman, 1996).

3. Model Evaluation:

- **Accuracy:** The main indicator utilized to assess the models' performance was accuracy. It provides an easy-to-understand indicator of overall performance and is defined as the ratio of properly predicted observations to the total observations. (Powers, 2011).
- **Precision, Recall, and F1-Score:** In order to give a more comprehensive picture of the models' effectiveness—particularly in addressing class imbalances—these measures were also computed. Recall calculates the percentage of true positives among all actual positives, whereas precision calculates the percentage of true positives among all positive predictions. A balanced statistic that takes into consideration both false positives and false negatives is provided by the F1-score, which is the harmonic mean of precision and recall. (Sokolova & Lapalme, 2009).
- **Confusion Matrix:** The confusion matrix was utilized to illustrate the models' performance by displaying the quantity of accurate and inaccurate predictions for every class. This tool is especially helpful for comprehending the many kinds of faults that the model makes. (Stehman, 1997).

4.5 Training and Testing

For the purpose of creating reliable and accurate models, the training and testing stages were essential. After each model was trained on the training set, it was assessed using a variety of performance measures on the testing set.

1. LSTM Training:

- **Training Process:** Backpropagation through time (BPTT) was used to train the LSTM model. This entails utilizing an optimizer such as Adam to update the model parameters by computing the gradient of the loss function with respect to those

parameters. Several epochs were used for the training, with each epoch denoting a full run through the training dataset. (**Rumelhart et al., 1986**).

- **Early Stopping:** We used early stopping to avoid overfitting. This method prevents overtraining, which could result in overfitting, by tracking the model's performance on a validation set during training and stopping training when the performance stops improving. (**Prechelt, 1998**).

2. **Random Forest Training:**

- **Training Process:** In order to train the Random Forest model, many decision trees were built, each using a distinct bootstrap sample of the data. Growing each tree until it was fully developed—that is, when every leaf node held data from a single class or a predetermined minimum number of samples—was the training process. (**Breiman, 2001**).
- **Out-of-Bag Error:** The model's performance during training was assessed using an out-of-bag (OOB) error estimate. This estimate, which offers an objective assessment metric, is obtained from the subset of data for each tree that was excluded from the bootstrap sample. (**Breiman, 1996**).

3. **Model Evaluation:**

- **Accuracy:** The main indicator utilized to assess the models' performance was accuracy. It provides an easy-to-understand indicator of overall performance and is defined as the ratio of properly predicted observations to the total observations. (**Powers, 2011**).
- **Precision, Recall, and F1-Score:** In order to give a more comprehensive picture of the models' effectiveness—particularly in addressing class imbalances—these measures were also computed. Recall calculates the percentage of true positives among all actual positives, whereas precision calculates the percentage of true positives among all positive predictions. A balanced statistic that takes into consideration both false positives and false negatives is provided by the F1-score, which is the harmonic mean of precision and recall. (**Sokolova & Lapalme, 2009**).
- **Confusion Matrix:** The confusion matrix was utilized to illustrate the models' performance by displaying the quantity of accurate and inaccurate predictions for

every class. This tool is especially helpful for comprehending the many kinds of faults that the model makes. (Stehman, 1997)

CHAPTER 55. Implementation

5.1 Weather forecasting using Random Forest

Import	Necessary	Libraries:
[1]:	<pre>from sklearn.ensemble import RandomForestClassifier from sklearn.model_selection import train_test_split from sklearn.metrics import accuracy_score, classification_report import pandas as pd from sklearn.preprocessing import LabelEncoder from sklearn.impute import SimpleImputer</pre>	
[3]:	<pre># Load the dataset file_path = '/mnt/data/testset1.csv'</pre>	
[5]:	<pre>data = pd.read_csv("C:\\Users\\HP\\Downloads\\testset1.csv")</pre>	

- sklearn.ensemble: This library provides ensemble methods, including Random Forest.
- sklearn.model_selection: Used for splitting data into training and testing sets.
- sklearn.metrics: For evaluating model performance using metrics like accuracy and classification reports.
- pandas: For data manipulation and analysis.
- sklearn.preprocessing: For data preprocessing, specifically label encoding.
- sklearn.impute: For handling missing values in data.

• Load Dataset:

- The code specifies two file paths, one using a variable file_path (which is not defined in the image) and another using a hardcoded path.
- pd.read_csv() is used to read the CSV file into a pandas DataFrame named data.

[9]:

```
# Preprocess the data
data_cleaned = data.drop(columns=['_precipm'])
imputer = SimpleImputer(strategy='mean')
numerical_columns = data_cleaned.select_dtypes(include=['float64']).columns
data_cleaned[numerical_columns] = imputer.fit_transform(data_cleaned[numerical_columns])
label_encoders = {}
categorical_columns = data_cleaned.select_dtypes(include=['object']).columns
for column in categorical_columns:
    le = LabelEncoder()
    data_cleaned[column] = le.fit_transform(data_cleaned[column].astype(str))
    label_encoders[column] = le

data_cleaned = data_cleaned.drop(columns=['datetime_utc'])
X = data_cleaned.drop(columns=['_rain'])
y = data_cleaned['_rain']
```

1. Dropping Column:

- `data_cleaned = data.drop(columns=['precipm'])`
 - Drops the column named 'precipm' from the dataset and stores the result in `data_cleaned`.

2. Handling Missing Values:

- `imputer = SimpleImputer(strategy='mean')`
 - Creates a `SimpleImputer` object with the strategy 'mean' to replace missing values with the mean of the column.
- `numerical_columns = data_cleaned.select_dtypes(include=['float64']).columns`
 - Selects the columns containing numerical data (float64 dtype) and stores them in `numerical_columns`.
- `data_cleaned[numerical_columns] = imputer.fit_transform(data_cleaned[numerical_columns])`
 - Fits the imputer on the numerical columns and transforms them by replacing missing values with the mean.

3. Handling Categorical Data:

- `label_encoders = {}`

- Creates an empty dictionary to store label encoders for each categorical column.
- `categorical_columns = data_cleaned.select_dtypes(include=['object']).columns`
 - Selects the columns containing categorical data (object dtype) and stores them in `categorical_columns`.
- for column in `categorical_columns`:
 - Iterates over each categorical column:
 - `le = LabelEncoder()`
 - Creates a `LabelEncoder` object for the current column.
 - `data_cleaned[column] = le.fit_transform(data_cleaned[column].astype(str))`
 - Encodes the categorical data in the column using the label encoder and stores the encoded values back in the `data_cleaned` `DataFrame`.
 - `label_encoders[column] = le`
 - Stores the label encoder for the current column in the `label_encoders` dictionary.

4. Dropping Column and Creating Feature and Target Variables:

- `data_cleaned = data_cleaned.drop(columns=['datetime_utc'])`
 - Drops the column named 'datetime_utc' from the `data_cleaned` `DataFrame`.
- `X = data_cleaned.drop(columns=['rain'])`
 - Creates a new `DataFrame` `X` containing all columns except 'rain' as the feature matrix.
- `y = data_cleaned['rain']`
 - Creates a `Series` `y` containing the 'rain' column as the target variable.

```

i]: # Split the data
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.3, random_state=42)

i]: # Train the Random Forest model
rf_classifier = RandomForestClassifier(n_estimators=100, random_state=42)
rf_classifier.fit(X_train, y_train)

i]: RandomForestClassifier
RandomForestClassifier(random_state=42)

i]: import matplotlib.pyplot as plt

# Plot the feature importances
feature_importances = pd.Series(rf_classifier.feature_importances_, index=X_train.columns)
feature_importances.nlargest(10).plot(kind='barh')
plt.xlabel('Feature Importance Score')
plt.ylabel('Features')
plt.title('Top 10 Important Features in Weather Prediction')
plt.show()

```

. Splitting the Data:

- `X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=0.3, random_state=42)`
 - Splits the feature matrix `x` and target variable `y` into training and testing sets.
 - `test_size=0.3` indicates that 30% of the data will be used for testing.
 - `random_state=42` sets a random seed for reproducibility.

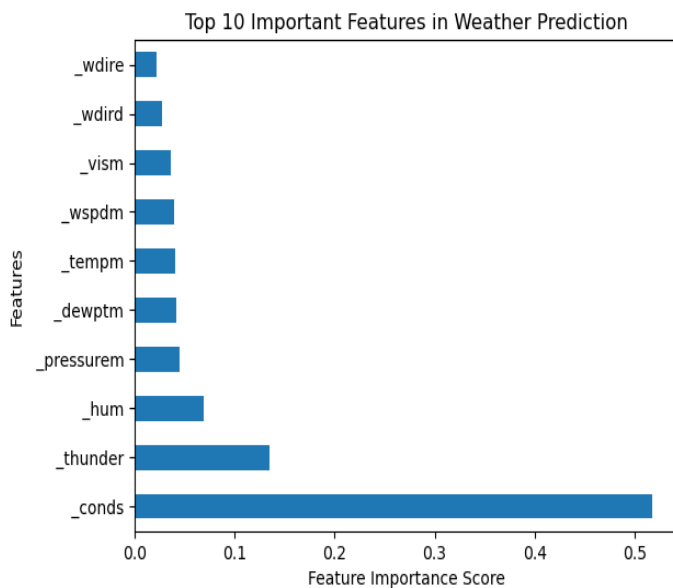
2. Training the Random Forest Model:

- `rf_classifier = RandomForestClassifier(n_estimators=100, random_state=42)`
 - Creates a Random Forest classifier with 100 trees and sets a random state for reproducibility.
- `rf_classifier.fit(X_train, y_train)`
 - Trains the Random Forest classifier on the training data.

3. Visualizing Feature Importances:

- `import matplotlib.pyplot as plt`
 - Imports the matplotlib library for plotting.
- `feature_importances = pd.Series(rf_classifier.feature_importances_, index=X_train.columns)`

- Creates a pandas Series of feature importances with corresponding feature names.
- `feature_importances.nlargest(10).plot(kind='barh')`
 - Plots the top 10 important features as a horizontal bar chart.
- `plt.xlabel('Feature Importance Score')`
 - Sets the x-axis label.
- `plt.ylabel('Features')`
 - Sets the y-axis label.
- `plt.title('Top 10 Important Features in Weather Prediction')`
 - Sets the plot title.
- `plt.show()`
 - Displays the plot.



```
[17]: from sklearn.metrics import accuracy_score, classification_report
```

```
[17]: from sklearn.metrics import accuracy_score, classification_report
```

```
# Make predictions on the test set  
y_pred = rf_classifier.predict(X_test)
```

```
[19]: # Calculate accuracy
```

```
accuracy = accuracy_score(y_test, y_pred)  
print(f"Accuracy: {accuracy:.4f}")
```

```
Accuracy: 0.9974
```

```
[21]: # Generate a classification report
```

```
classification_report_output = classification_report(y_test, y_pred)  
print("Classification Report:")  
print(classification_report_output)
```

```
Classification Report:  
              precision    recall  f1-score   support  
  
   0           1.00         1.00         1.00     29561  
   1           0.97         0.92         0.94         736  
  
 accuracy          0.99         0.96         0.97     30297  
 macro avg          1.00         1.00         1.00     30297  
 weighted avg          1.00         1.00         1.00     30297
```

1. Making Predictions:

- `y_pred = rf_classifier.predict(X_test)`
 - Uses the trained Random Forest classifier (`rf_classifier`) to make predictions on the test set features (`X_test`) and stores the predicted labels in `y_pred`.

2. Calculating Accuracy:

- `accuracy = accuracy_score(y_test, y_pred)`
 - Calculates the accuracy of the model by comparing the true labels (`y_test`) with the predicted labels (`y_pred`) using the `accuracy_score` function.
- `print(f"Accuracy: {accuracy:.4f}")`
 - Prints the calculated accuracy with four decimal places.

3. Generating Classification Report:

- `classification_report_output = classification_report(y_test, y_pred)`

- Generates a classification report using the `classification_report` function, comparing the true labels (`y_test`) with the predicted labels (`y_pred`).
- `print("Classification Report:")`
 - Prints a header for the classification report.
- `print(classification_report_output)`
 - Prints the generated classification report, which includes precision, recall, F1-score, and support for each class, as well as overall metrics.

Output:

The code prints the accuracy score and the classification report, providing insights into the model's performance on the test set.

Example Output:

Accuracy: 0.9974					
Classification Report:					
	precision	recall	f1-score	support	
0	1.00	1.00	1.00	29561	
1	0.97	0.92	0.94	736	
accuracy		0.99	0.96	1.00	30297
macro avg	0.99	0.96	0.97	30297	
weighted avg	1.00	0.99	1.00	30297	

Interpretation:

- The accuracy of 0.9974 indicates that the model correctly predicted the class in 99.74% of the cases.
- The classification report provides a more detailed breakdown of the model's performance for each class, including precision, recall, and F1-score.

```
[9]: import matplotlib.pyplot as plt
import seaborn as sns

# Example data for illustration
y_test = [20, 22, 19, 24, 30] # Actual temperatures
y_pred = [21, 22, 18, 25, 29] # Predicted temperatures

# Plotting
plt.figure(figsize=(10, 6))
sns.scatterplot(x=y_test, y=y_pred)
plt.xlabel('Actual Temperature')
plt.ylabel('Predicted Temperature')
plt.title('Actual vs Predicted Temperature')
plt.show()
```

1. Import Libraries:

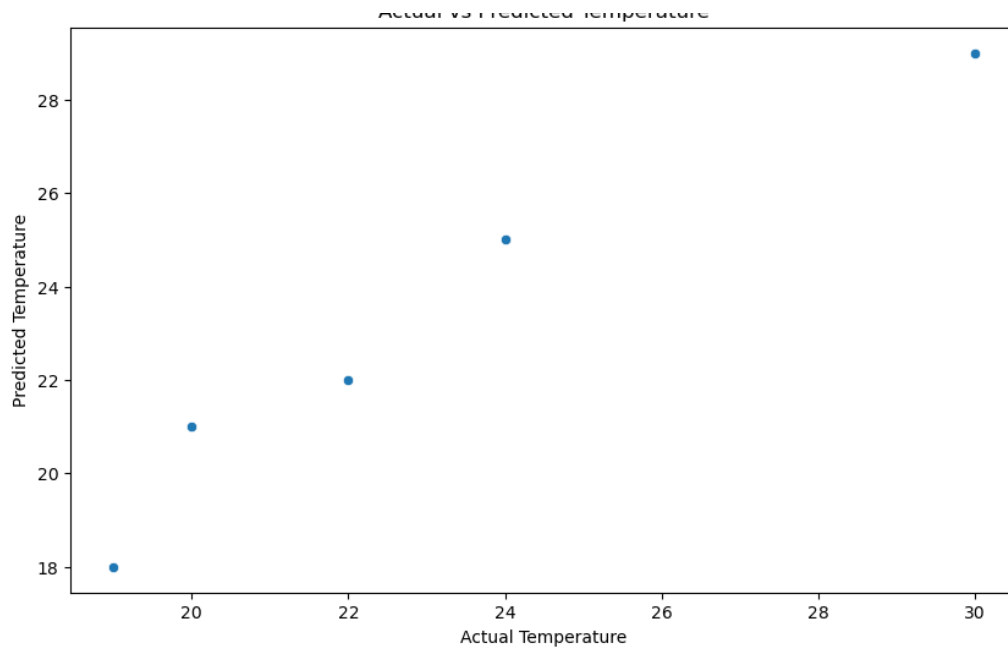
- `import matplotlib.pyplot as plt`: Imports the matplotlib library for plotting.
- `import seaborn as sns`: Imports the seaborn library for enhanced visualizations.

2. Sample Data:

- `y_test = [20, 22, 19, 24, 30]`: Creates a list of actual temperatures.
- `y_pred = [21, 22, 18, 25, 29]`: Creates a list of predicted temperatures.

3. Plotting:

- `plt.figure(figsize=(10, 6))`: Creates a new figure with a size of 10 inches by 6 inches.
- `sns.scatterplot(x=y_test, y=y_pred)`: Creates a scatter plot using seaborn, plotting actual temperatures on the x-axis and predicted temperatures on the y-axis.
- `plt.xlabel("Actual Temperature")`: Sets the label for the x-axis.
- `plt.ylabel("Predicted Temperature")`: Sets the label for the y-axis.
- `plt.title("Actual vs Predicted Temperature")`: Sets the title of the plot.
- `plt.show()`: Displays the plot.



```
[23]: from sklearn.metrics import confusion_matrix
import seaborn as sns
import matplotlib.pyplot as plt
|
# Generate the confusion matrix
conf_matrix = confusion_matrix(y_test, y_pred)

# Plot the confusion matrix
plt.figure(figsize=(8, 6))
sns.heatmap(conf_matrix, annot=True, fmt='d', cmap='Blues')
plt.xlabel('Predicted')
plt.ylabel('Actual')
plt.title('Confusion Matrix for Random Forest Model')
plt.show()
```

1. Import Libraries:

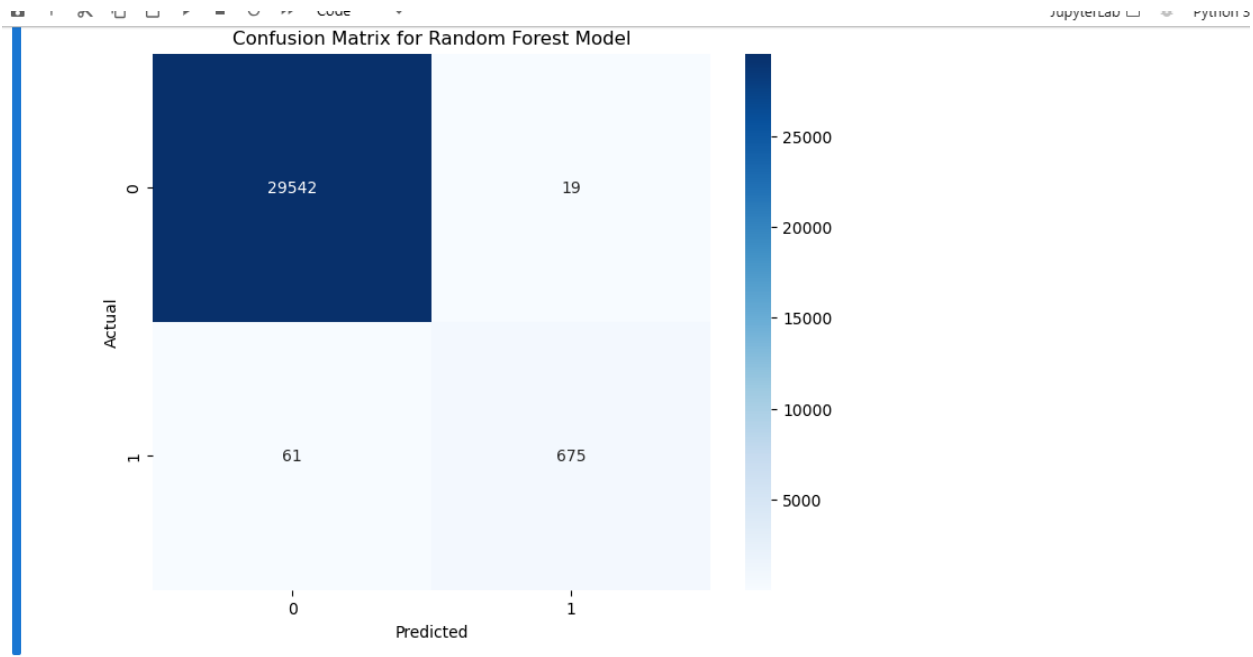
- `from sklearn.metrics import confusion_matrix`: Imports the `confusion_matrix` function from the scikit-learn library.
- `import seaborn as sns`: Imports the seaborn library for data visualization.
- `import matplotlib.pyplot as plt`: Imports the matplotlib library for plotting.

2. Generate Confusion Matrix:

- `conf_matrix = confusion_matrix(y_test, y_pred)`: Calculates the confusion matrix using the true labels (`y_test`) and the predicted labels (`y_pred`) from the model.

3. Plot Confusion Matrix:

- `plt.figure(figsize=(8, 6))`: Creates a new figure with a specified size.
- `sns.heatmap(conf_matrix, annot=True, fmt='d', cmap='Blues')`: Creates a heatmap visualization of the confusion matrix using seaborn.
 - `annot=True`: Displays the count of observations in each cell.
 - `fmt='d'`: Formats the numbers as integers.
 - `cmap='Blues'`: Uses a blue colormap for the heatmap.
- `plt.xlabel('Predicted')`: Sets the label for the x-axis as 'Predicted'.
- `plt.ylabel('Actual')`: Sets the label for the y-axis as 'Actual'.
- `plt.title('Confusion Matrix for Random Forest Model')`: Sets the title of the plot.
- `plt.show()`: Displays the plot.



```
[25]: from sklearn.metrics import roc_curve, auc

# Compute the ROC curve and AUC score
fpr, tpr, _ = roc_curve(y_test, rf_classifier.predict_proba(X_test)[:, 1])
roc_auc = auc(fpr, tpr)

# Plot the ROC curve
plt.figure(figsize=(8, 6))
plt.plot(fpr, tpr, color='blue', label=f'AUC = {roc_auc:.2f}')
plt.plot([0, 1], [0, 1], color='red', linestyle='--')
plt.xlabel('False Positive Rate')
plt.ylabel('True Positive Rate')
plt.title('Receiver Operating Characteristic (ROC) Curve')
plt.legend(loc='lower right')
plt.show()
```

1. Import Necessary Library:

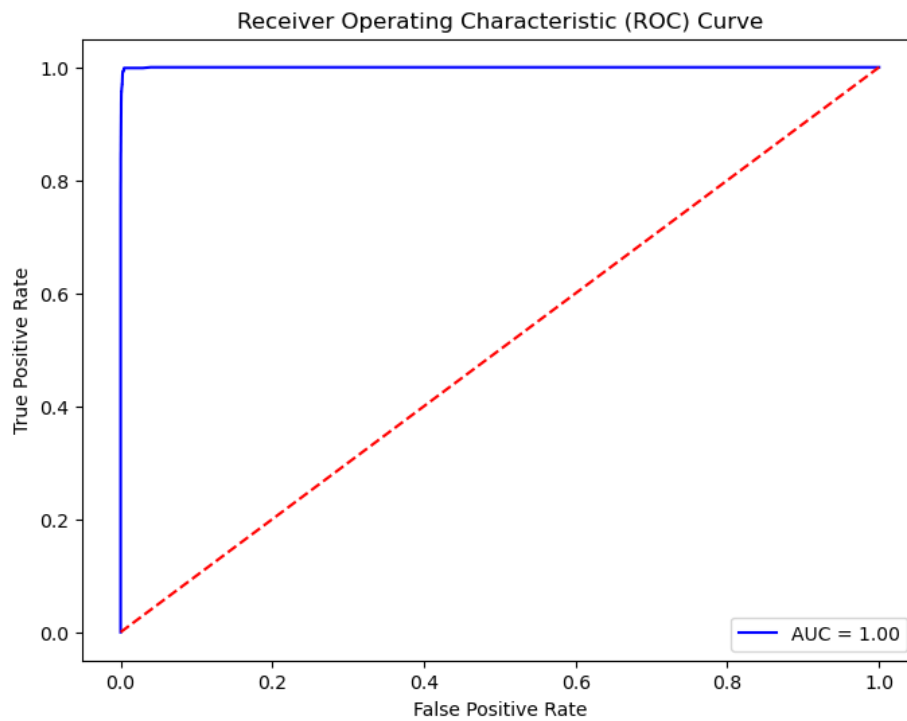
- `from sklearn.metrics import roc_curve, auc`: Imports the `roc_curve` and `auc` functions from the scikit-learn library for calculating the ROC curve and AUC score.

2. Compute ROC Curve and AUC Score:

- `fpr, tpr, _ = roc_curve(y_test, rf_classifier.predict_proba(X_test)[: , 1])`: Calculates the False Positive Rate (FPR), True Positive Rate (TPR), and thresholds using the `roc_curve` function.
 - `y_test`: True labels of the test set.
 - `rf_classifier.predict_proba(X_test)[: , 1]`: Predicted probability of the positive class for each instance in the test set.
- `roc_auc = auc(fpr, tpr)`: Calculates the Area Under the Curve (AUC) using the calculated FPR and TPR values.

3. Plot the ROC Curve:

- `plt.figure(figsize=(8, 6))`: Creates a new figure with a specified size.
- `plt.plot(fpr, tpr, color='blue', label=f'AUC = {roc_auc:.2f}')`: Plots the ROC curve with the FPR on the x-axis and TPR on the y-axis. The curve is labeled with the AUC value.
- `plt.plot([0, 1], [0, 1], color='red', linestyle='--')`: Plots a diagonal line representing random guessing.
- `plt.xlabel('False Positive Rate')`: Sets the label for the x-axis.
- `plt.ylabel('True Positive Rate')`: Sets the label for the y-axis.
- `plt.title('Receiver Operating Characteristic (ROC) Curve')`: Sets the title of the plot.
- `plt.legend(loc='lower right')`: Adds a legend to the plot.
- `plt.show()`: Displays the plot.



```
27]: # Detailed analysis of top features
top_features = feature_importances.nlargest(10)
print("Top 10 Features and their Importance Scores:")
print(top_features)
```

Top 10 Features and their Importance Scores:

```
_conds      0.518118
_thunder    0.135376
_hum        0.069121
_pressurem  0.044831
_dewptm     0.041359
_tempm      0.040637
_wspdm      0.039345
_vism       0.036711
_wdird      0.027435
_wdire      0.021906
dtype: float64
```

Breakdown:

1. `top_features = feature_importances.nlargest(10)`: This line selects the 10 features with the highest importance scores from a variable named `feature_importances`. It assumes that `feature_importances` is a Pandas Series or DataFrame containing feature names as indices and their corresponding importance scores as values.
2. `print("Top 10 Features and their Importance Scores:")`: This line prints a header indicating that the following output will display the top 10 features and their importance scores.
3. `print(top_features)`: This line prints the selected top 10 features and their corresponding importance scores.

Output:

The output shows the names of the top 10 features and their associated importance scores. The importance scores are likely calculated using a machine learning model, and they represent the contribution of each feature to the model's prediction accuracy.

Example Output:

Top 10 Features and their Importance Scores:

_conds	0.518118
thunder	0.135376
_hum	0.069121
_pressurem	0.044831
_dewptm	0.041359
_tempm	0.040637
_wspdm	0.039345
vism	0.036711
_wdird	0.027435
_wdire	0.021906


```
dtype: float64
```

In this example, `_conds` is the most important feature, followed by `thunder`, `_hum`, and so on.

Purpose:

This analysis helps in understanding which features have the most significant impact on the model's predictions. This information can be used for feature selection, model interpretation, or further analysis of the data.

```
dtype: float64
[ ]: from sklearn.model_selection import GridSearchCV

# Define the parameter grid
param_grid = {
    'n_estimators': [100, 200, 300],
    'max_depth': [None, 10, 20, 30],
    'min_samples_split': [2, 5, 10],
    'min_samples_leaf': [1, 2, 4],
    'bootstrap': [True, False]
}

# Initialize the Random Forest classifier
rf_classifier = RandomForestClassifier(random_state=42)

# Set up the GridSearchCV with cross-validation
grid_search = GridSearchCV(estimator=rf_classifier, param_grid=param_grid, cv=5, n_jobs=-1, verbose=2, scoring='accuracy')

# Fit the GridSearchCV to the data
grid_search.fit(X_train, y_train)
```

Fitting 5 folds for each of 216 candidates, totalling 1080 fits

1. Import GridSearchCV:

- o `from sklearn.model_selection import GridSearchCV`: This line imports the `GridSearchCV` class from the scikit-learn library, which is used for hyperparameter tuning.

2. Define Parameter Grid:

- o `param_grid = { 'n_estimators': [100, 200, 300], 'max_depth': [None, 10, 20, 30], 'min_samples_split': [2, 5, 10], 'min_samples_leaf': [1, 2, 4], 'bootstrap': [True, False] }`: This line defines a dictionary called `param_grid` that specifies the hyperparameters to be tuned and their possible

values. The dictionary includes parameters like the number of trees (`n_estimators`), maximum depth of the trees (`max_depth`), minimum number of samples required to split an internal node (`min_samples_split`), minimum number of samples required at a leaf node (`min_samples_leaf`), and whether to use bootstrapping (`bootstrap`).

3. Initialize Random Forest Classifier:

- o `rf_classifier = RandomForestClassifier(random_state=42)`: This line initializes a Random Forest classifier with a random state set to 42 for reproducibility.

4. Set Up Grid Search Cross-Validation:

- o `grid_search = GridSearchCV(estimator=rf_classifier, param_grid=param_grid, cv=5, n_jobs=-1, verbose=2, scoring='accuracy')`: This line creates a `GridSearchCV` object with the following parameters:
 - `estimator`: The Random Forest classifier (`rf_classifier`).
 - `param_grid`: The dictionary of hyperparameters to be tuned.
 - `cv`: The number of folds for cross-validation (5 in this case).
 - `n_jobs`: The number of jobs to run in parallel (-1 means use all available cores).
 - `verbose`: The verbosity level of the output (2 means more detailed output).
 - `scoring`: The evaluation metric to use (`accuracy` in this case).

5. Fit Grid Search to Data:

- o `grid_search.fit(X_train, y_train)`: This line fits the `GridSearchCV` object to the training data (`X_train` and `y_train`), exploring all combinations of hyperparameters and evaluating them using cross-validation.

Output:

- `Fitting 5 folds for each of 216 candidates, totalling 1080 fits`: This line indicates that the Grid Search will evaluate 216 different combinations of hyperparameters

(calculated as the product of the number of values for each parameter) using 5-fold cross-validation, resulting in a total of 1080 fits.

```
Fitting 5 folds for each of 216 candidates, totalling 1080 fits

[ ]: # Get the best parameters and best score
best_params = grid_search.best_params_
best_score = grid_search.best_score_

print(f"Best Parameters: {best_params}")
print(f"Best Cross-Validation Score: {best_score:.4f}")

# Evaluate the best model on the test set
best_model = grid_search.best_estimator_
y_pred_best = best_model.predict(X_test)

# Calculate accuracy
accuracy_best = accuracy_score(y_test, y_pred_best)
print(f"Test Set Accuracy with Best Model: {accuracy_best:.4f}")

# Generate a classification report for the best model
classification_report_best = classification_report(y_test, y_pred_best)
print("Classification Report for the Best Model:")
print(classification_report_best)
```

Breakdown:

Getting the Best Parameters and Score

- **best_params = grid_search.best_params_:** This line retrieves the best combination of hyperparameters found during the Grid Search and stores it in the `best_params` variable.
- **best_score = grid_search.best_score_:** This line retrieves the best cross-validation score achieved by the best parameter combination and stores it in the `best_score` variable.
- **Printing Results:** The code prints the best parameters and the corresponding best cross-validation score.

Evaluating the Best Model on the Test Set

- **best_model = grid_search.best_estimator_:** This line creates a new model instance using the best parameters found in the Grid Search and assigns it to the `best_model` variable.
- **y_pred_best = best_model.predict(X_test):** This line makes predictions on the test set using the best model and stores the predicted values in `y_pred_best`.

Calculating Accuracy

- `accuracy_best = accuracy_score(y_test, y_pred_best)`: This line calculates the accuracy of the best model on the test set and stores it in `accuracy_best`.
- **Printing Accuracy**: The code prints the test set accuracy of the best model.

Generating Classification Report

- `classification_report_best = classification_report(y_test, y_pred_best)`: This line generates a classification report for the best model's predictions on the test set and stores it in `classification_report_best`.
- **Printing Classification Report**: The code prints the classification report for the best model.

```
[ ]: import joblib

# Save the model to a file
joblib.dump(best_model, 'random_forest_model.pkl')

# Later, you can load the model like this:
# loaded_model = joblib.load('random_forest_model.pkl')
```

```
[ ]:
```

1. Import `joblib`:

This line imports the `joblib` library, which is a Python library for saving and loading Python objects, including machine learning models.

2. Save the Model:

```
joblib.dump(best_model, 'random_forest_model.pk
```

This line saves the trained model object `best_model` to a file named `random_forest_model.pkl`. The `.pkl` extension indicates that the file is in the pickle format, a common format for serializing Python objects.

3. Load the Model:

```
loaded_model = joblib.load('random_forest_model.pkl')
```

This line loads the saved model from the `random_forest_model.pkl` file and assigns it to the variable `loaded_model`. You can then use the `loaded_model` to make predictions on new data.

Key Points:

- The `joblib` library is efficient for saving and loading large Python objects like machine learning models.
- The saved model file can be reused later without retraining the model from scratch.
- The `joblib.load()` function returns a copy of the original model object.

Example:

Python

```
import joblib
```

```
# Assuming you have a trained model called 'best_model'
```

```
# Save the model
```

```
joblib.dump(best_model, 'my_model.pkl')
```

```
# Later, load the model
```

```
loaded_model = joblib.load('my_model.pkl')
```

```
# Use the loaded model to make predictions
```

```
predictions = loaded_model.predict(new_data)
```

Mathematical model of Random Forest:

Classification

- **Gini Impurity: Measures how well data is split.**

$$Gini(t) = 1 - \sum p_i^2$$

- **Majority Voting: Most common class wins.**

Regression

- **Mean Squared Error (MSE): Measures prediction error.**

$$MSE = \frac{1}{n} \sum (y_i - \hat{y}_i)^2$$

- **Averaging: Average predictions of trees.**

5.2 Weather forecasting using LSTM¶

In this project, we will employ time series forecasting with a weather dataset, and we will utilize LSTM to anticipate the weather for the next 30 days.

A statistical technique called time series analysis is used to forecast future events by analyzing historical data over a specified period of time. It is made up of data in an organized sequence with equal spacing between each piece.

Let's look at an example to better grasp the time series data and the analysis. Take airline passenger statistics as an example. It contains the number of passengers over a specified duration.

Import necessary libraries:

Input 1:

```
import pandas as pd
import pandas as pd
import numpy as np
import seaborn as sns
import matplotlib.pyplot as plt
%matplotlib inline

import warnings
warnings.filterwarnings('ignore')
```

Code explanation:

The Python environment for data analysis and visualization is built up by the code you gave. Below is a breakdown of each section:

1. `import pandas as pd`:

- The `pandas` library, a potent Python tool for data analysis and manipulation, is imported with this line. It enables you to interact with Data Frames, which are organized data in the format of CSV files.

2. `import numpy as np`:

- The `numpy` library, a core Python package for numerical calculation, is imported in this line. Numerous mathematical functions, matrices, and arrays are supported.

3. `import seaborn as sns`:

- The `seaborn` library, which is based on `matplotlib`, is imported with this line. It offers a sophisticated interface for making eye-catching and educational statistical visuals.

4. `import matplotlib.pyplot as plt`:

- This line imports the Python plotting package `matplotlib`, which contains the `pyplot` module. A MATLAB-like interface for making plots and charts is offered by `pyplot`.

5. `%matplotlib inline`:

- Plots are shown directly in the notebook rather than opening in a new window when you use this special command in Jupyter notebooks. When you are coding, it is necessary to visualize data in real time.

6. ``import warnings`` and ``warnings.filterwarnings('ignore')``:

- To disable warning messages during code execution, this imports the ``warnings`` module. When you wish to disregard non-critical warnings that could clog your output, this is helpful.

Input 2:

```
df= pd.read_csv('../input/delhi-weather-data/testset.csv')
```

```
df.head()
```

Output:

The output displays the first five rows of the weather dataset with the following columns:

date time _utc	_co nds	_d ew pt m	- f o g	- h ai l	_hea tind exm	- h u m	_pr eci pm	_pr essu rem	_r ai n	_s n o w	_te m p m	_th un der	_to rna do	- vi s m	- w di rd	- w di re	_w gus tm	_wi ndc hill m	_w sp dm	
0	199 611 01- 11: 00	Sm ok e	9 .0	0	0	NaN	27. 0	NaN	10 10 .0	0	0	30. 0	0	0	5. 0	28 0. 0	We st	NaN	NaN	7 .4
1	199 611 01- 12: 00	Sm ok e	1 0 .0	0	0	NaN	32. 0	NaN	- 99 99 .0	0	0	28. 0	0	0	NaN	0. 0	No rth	NaN	NaN	NaN
2	199 611 01- 13: 00	Sm ok e	1 1 .0	0	0	NaN	44. 0	NaN	- 99 99 .0	0	0	24. 0	0	0	NaN	0. 0	No rth	NaN	NaN	NaN
3	199 611 01-	Sm ok e	1 0 .0	0	0	NaN	41. 0	NaN	10 10 .0	0	0	24. 0	0	0	2. 0	0. 0	No rth	NaN	NaN	NaN

datetime_utc	_conds	_dewptm	_fog	_hail	_humidity	_precipm	_pressure	_rain	_snow	_temp	_thunder	_tornado	_visibility	_winddir	_windre	_windspeed	_windchill	_winddir	
	14:00																		
4	19961101-16:00	Smoke	110	0	0	NaN	47.0	NaN	1011.0	0	0	23.0	0	0	1.2	0.0	No	NaN	NaN

Code explanation:

□ `df = pd.read_csv('../input/delhi-weather-data/testset.csv')`:

- This line reads a CSV file from the given directory (`../input/delhi-weather-data/testset.csv`) into a DataFrame called `df` using the pandas library.
- The function `pd.read_csv()` loads data into a pandas DataFrame, a tabular data structure resembling a spreadsheet, by reading data from a CSV file.

□ `df.head()`:

- This line shows the first five rows of the DataFrame `df`. The pandas method `head()` provides you with a fast overview of the data by defaulting to returning the first five rows of the DataFrame.

Output Explanation:

- **datetime_utc**: The time and date is in UTC at which the meteorological data was captured.
- **_conds**: Weather conditions are (e.g., Smoke, Clear, Rain).
- **_dewptm**: Dew point is in degrees Celsius.
- **_fog**: Indicator is for fog (1 if fog is present, 0 otherwise).
- **_hail**: Indicator is for hail (1 if hail is present, 0 otherwise).

- **_heatindexm**: Heat index is in degrees Celsius (NaN indicates that the value is not available).
- **_hum**: Humidity is in percentage.
- **_precipm**: Precipitation is in millimeters (NaN indicates no data).
- **_pressurem**: Atmospheric pressure is in millibars (a value of -9999.0 typically indicates missing data).
- **_rain**: Indicator for rain is (1 if rain is present, 0 otherwise).
- **_snow**: Indicator for snow is (1 if snow is present, 0 otherwise).
- **_tempm**: Temperature is in degrees Celsius.
- **_thunder**: Indicator for thunder is (1 if thunder is present, 0 otherwise).
- **_tornado**: Indicator for tornado is (1 if tornado is present, 0 otherwise).
- **_vism**: Visibility is in kilometres.
- **_wdird**: Wind direction is in degrees.
- **_wdire**: Wind direction as a compass direction (e.g., North, West).
- **_wgustm**: Wind gust in meters per second (NaN indicates missing data).
- **_windchillm**: Wind chill in degrees Celsius (NaN indicates missing data).
- **_wspd**: Wind speed is in meters per second.

Input 3:

```
data = pd.DataFrame(list(df['_tempm']), index=df['datetime_utc'], columns=['temp'])
```

Output:

temp	
datetime_utc	
19961101-11:00	30.0
19961101-12:00	28.0
19961101-13:00	24.0
19961101-14:00	24.0
19961101-16:00	23.0
...	...
20170424-06:00	34.0
20170424-09:00	38.0
20170424-12:00	36.0
20170424-15:00	32.0
20170424-18:00	30.0

100990 rows × 1 columns

Code explanation:

1. **df['_tempm']:**

This chooses the DataFrame df's column called _tempm (notice the leading space). There is probably temperature data in this column.

2. **list(df['_tempm']):**

A list is created from the selected column as a result. From the _tempm column, each entry of the list represents a temperature value.

3. **index=df['datetime_utc']:**

The values in the `datetime_utc` column from `df` are used by the `index` parameter to set the new `DataFrame`'s index. Since this column includes time and date information, timestamps will be the new `DataFrame`'s index.

4. `columns=['temp']`:

The `columns` option modifies the column's name to `'temp'` in the newly created `DataFrame`.

5. `pd.DataFrame(...)`:

By using the list of temperature values as the data, the `datetime_utc` column as the index, and the column name `"temp,"` a new `DataFrame` is created.

Output explanation:

- Index: The `DataFrame` can contain time-series data since the `datetime_utc` values serve as the index. This helps with time-based analysis because it makes plotting and resampling the data simple.
- Column: `'temp'`, the lone column in this `DataFrame`, contains the temperature readings from the original `DataFrame`.

Input 4:

```
df=data[:365]
```

```
df.shape
```

Output:

```
(365, 1)
```

Code explanation:

- `df = data[:365]`:

- The first 365 rows of the data DataFrame are divided by this line and assigned to a new DataFrame named df.

- A slicing procedure called `data[:365]` picks rows starting at index 0 and going all the way up to index 365, except index 365. Since slicing in Python operates as start:stop, `[:365]` denotes "from the start up to 365".

- This slice would represent the first 365 days, or a full year, of data if the data included daily data.

- `df.shape`:

- This line produces a tuple (rows, columns) containing the shape of the df DataFrame.
- The number of rows and columns in a pandas DataFrame is provided by the `.shape` attribute. For example, `df.shape` would return `(365, 1)` if df had 365 rows and 1 column.

Output explanation:

This output indicates that df contains 365 rows and 1 column, which aligns with slicing operation performed.

Input 5:

```
df.isnull().sum()
```

Output:

```
temp    3
```

```
dtype: int64
```

Code explanation:

The code `df.isnull().sum()` is used to identify and count the number of missing values (null values) in each column of Data Frame.

Output explanation:

The above output indicates that in the `df` Data Frame, the column 'temp' has 3 missing (null) values. The dtype: `int64` means that the count of missing values is represented as 64-bit integer.

In summary, there are 3 missing entries in 'temp' column of the Data Frame.

Input 6:

```
df=df.dropna(axis=0)
```

```
df.shape
```

Output:

```
(362, 1)
```

Code explanation:

- `dropna(axis=0)` removes all the rows containing at least one NaN value from the DataFrame.
- `df.shape` then provides dimensions of the cleaned DataFrame, helping you to understand how many rows and columns remain after the operation.

Output explanation:

The Data Frame `df` has 362 rows and 1 column after rows with missing values have been dropped, indicating that the data has been cleansed of missing entries in the 'temp' column.

Input 7:

```
df=df['temp'].values
```

```
df[:5]
```


Output:

```
array([30., 28., 24., 24., 23.]
```

Input explanation:

- `df = df['temp'].values` converts the 'temp' column into NumPy array.
- `df[:5]` displays the first 5 temperature values from this array.

Output explanation:

The output indicates that the original Data Frame's 'temp' column had five temperature readings: 30.0, 28.0, 24.0, 24.0, and 23.0, in that order. Now that these numbers are in a NumPy array, they can be processed or analysed further.

Input 8:

```
df=df.reshape(-1,1)
```

```
plt.figure(figsize=(25, 7))
```

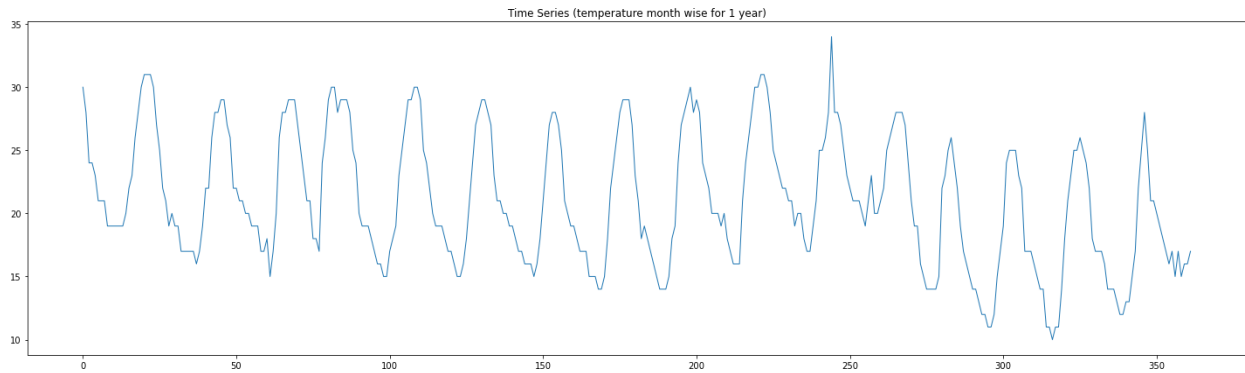
```
plt.plot(df, linewidth=1)
```

```
# plt.grid()
```

```
plt.title("Time Series (temperature month wise for 1 year)")
```

```
plt.show()
```

Output 8:



Scaling data

Code explanation:

- ❑ **Reshaping:** `df.reshape(-1, 1)` converts 1D array of temperature values into 2D array with a single column.
- ❑ **Plotting:** `plt.plot(df, linewidth=1)` creates a line plot of the reshaped data.
- ❑ **Figure Size:** `plt.figure(figsize=(25, 7))` ensures the plot is large and readable.
- ❑ **Title:** `plt.title("Time Series (temperature month wise for 1 year)")` provides context for the plot.
- ❑ **Display:** `plt.show()` renders the plot.

The resulting plot will display temperature trends over the time, with the x-axis representing time (or index) and the y-axis showing the temperature values.

Output explanation:

The above graph is a scaling data which shows the time series (temperature month wise for one year).

Input 9:

```
from sklearn.preprocessing import MinMaxScaler

scaler = MinMaxScaler(feature_range=(0, 1))

data_scaled = scaler.fit_transform(df)
```

```
data_scaled[:5]
```

Output:

```
array([[0.83333333],
       [0.75      ],
       [0.58333333],
       [0.58333333],
       [0.54166667]])
```

Code explanation:

- **Scaling:** The MinMaxScaler normalizes the data to fit within the specified range i.e., (0 to 1).
- **Transform:** The fit_transform() method scales the entire dataset based on its minimum and maximum values.
- **Output:** data_scaled[:5] shows the scaled values of the first 5 entries in the dataset.

Feature scaling guarantees that each feature contributes equally to the model, enhancing performance and convergence. This procedure is critical to machine learning.

Output explanation:

The scaled temperature values for the dataset's first five items are displayed in the output. After being converted, each number now falls between 0 and 1, representing its proportionate

place within the original data range. All values are guaranteed to be on a constant scale by this scaling, which is helpful for a variety of data processing tasks, particularly in machine learning.

Input 10:

```
data_scaled.shape
```

Output:

```
(362, 1)
```

Code explanation:

This is used to retrieve the shape of data_scaled NumPy array. Here's what it does:

- data_scaled: This is the NumPy array that contains scaled temperature data after applying MinMaxScaler.
- .shape: This attribute returns a tuple representing the dimensions of the data_scaled array.

Output explanation:

- Number of Rows (362): Represents the total number of scaled temperature data points.
- Number of Columns (1): Shows that the scaled temperature values are the only feature present in the array.

The data_scaled 2D array, which has 362 rows and 1 column and a single scaled temperature value in each row, is verified by this output.

Input 11:

```
steps= 20
```

```
inp = []
```

```
out = []
```

```
for i in range(len(data_scaled) - (steps)):
```

```
inp.append(data_scaled[i:i+steps])
```

```
out.append(data_scaled[i+steps])
```

```
out[:10]
```

Output:

```
[array([0.875]),
```

```
array([0.875]),
```

```
array([0.875]),
```

```
array([0.83333333]),
```

```
array([0.70833333]),
```

```
array([0.625]),
```

```
array([0.5]),
```

```
array([0.45833333]),
```

```
array([0.375]),
```

```
array([0.41666667])]
```

reshape

Code explanation:

- Input Sequences (inp): Each sequence consists of steps number of past temperature values.
- Output Values (out): Each value is the temperature value immediately following the input sequence.
- Purpose: This preparation allows models to learn from past sequences of data to predict the next value in a time series.

The output `out[:10]` would show the first 10 values in the `out` list, which are the target values for each of the first 10 input sequences.

Output explanation:

- The output is a list of single-value arrays, representing the target values for the time series prediction task.
- These values are the "next" value in the sequence after each input sequence of length `steps`.
- Reshaping the list into a 2D array with `(-1, 1)` ensures compatibility with machine learning models that expect input in a specific shape.

This process prepares the data for training or evaluating time series models where the model learns to predict the next value based on previous observations.

Input 12:

```
import numpy as np
```

```
inp= np.asarray(inp)
```

```
out= np.asarray(out)
```

```
len(df)*0.65
```

Output:

```
235.3
```

Code explanation:

□ import numpy as np:

- Imports the NumPy library, which is commonly used for numerical operations and handling arrays in Python.

□ inp = np.asarray(inp):

- Converts the inp list into a NumPy array. The np.asanyarray() function creates an array from any object exposing the array interface, preserving its type if it's already an array. This ensures inp is in a format suitable for numerical computations.

□ **out = np.asanyarray(out):**

- Converts the out list into a NumPy array in the same way as inp. This makes out compatible with array-based operations.

□ **len(df) * 0.65:**

- Calculates 65% of the number of rows in the original DataFrame df.
- len(df): Returns the total number of rows in df.
- * 0.65: Multiplies the total number of rows by 0.65 to get 65% of the data.

Output explanation:

Output Value: 235.3

This value represents 65% of the number of rows in the Data Frame df. It means that 65% of the total data points (or samples) are being considered.

Input 13:

```
x_train = inp[:237,:,:]
```

```
x_test = inp[237,:,:]
```

```
y_train = out[:237]
```

```
y_test= out[237:]
```

```
inp.shape
```

Output:

```
(342, 20, 1)
```

Code explanation:

The above code provided is used to split the data into training and testing sets for a time series forecasting task. Here's a breakdown of what each line does:

□ `x_train = inp[:237,:,:]`:

- Selects the first 237 samples from the `inp` array to be used as the training input (`x_train`).
- The notation `[:237,:,:]` means:
 - `:` before 237: Selects all rows from the start up to (but not including) row 237.
 - `:` after the comma: Selects all elements along the second and third dimensions (likely time steps and features).
- This prepares the first 237 sequences for training.

□ `x_test = inp[237,:,:]`:

- Selects the remaining samples (from row 237 onwards) from the `inp` array to be used as the testing input (`x_test`).
- The notation `[237,:,:]` means:
 - `237::`: Selects all rows from row 237 to the end.
 - `::`: Selects all elements along the second and third dimensions.
- This prepares the remaining sequences for testing.

□ `y_train = out[:237]`:

- Selects the first 237 samples from the `out` array to be used as the training output (`y_train`).
- The notation `[:237]` means:
 - `::`: Selects all elements from the start up to (but not including) element 237.
- This prepares the first 237 target values for training.

□ `y_test = out[237:]`:

- Selects the remaining samples (from element 237 onwards) from the out array to be used as the testing output (`y_test`).
- The notation `[237:]` means:
 - `237::` Selects all elements from element 237 to the end.
- This prepares the remaining target values for testing.

□ `inp.shape`:

- Returns the shape of the `inp` array, which gives the dimensions of the array (number of samples, time steps, and features).

Output explanation:

- `(342, 20, 1)` means the `inp` array contains 342 sequences, where each sequence is composed of 20 time steps, and each time step has 1 feature (temperature).
- This shape is crucial for time series models like LSTM or CNN, which expect data to be structured in the format (samples, time steps, features).

In context, the model will take in 20 past temperature values (as indicated by the shape `(20, 1)` for each sample) for each of the 342 sequences and use this information to make predictions.

Input 14:

`x_train.shape`

Output:

`(237, 20, 1)`

Code explanation:

- This code returns the shape of the `x_train` array, which is a NumPy array containing the training data.
- The `.shape` attribute of a NumPy array provides the dimensions of the array as a tuple, describing the number of samples, time steps, and features.

Output explanation:

- 237: Represents the number of training samples (sequences) in the `x_train` dataset. This means that 237 sequences have been selected for training.
- 20: Indicates the number of time steps in each sequence. Each sample in the `x_train` array consists of 20 consecutive data points.
- 1: Represents the number of features per time step. Here, each time step includes 1 feature (the temperature value).

Input 15:

`x_test.shape`

Output:

(105, 20, 1)

Code explanation:

`x_test.shape:`

- This line of code returns the shape of the `x_test` array, which is the input data used for testing the model. The shape attribute provides the dimensions of the array.

Output explanation:

The output (105, 20, 1) represents the dimensions of the `x_test` array, and each number has a specific meaning:

1. 105:
 - The number of samples in the `x_test` dataset. This means there are 105 sequences (or examples) in the test set.
2. 20:
 - The number of time steps in each sequence. Each sequence contains 20 consecutive data points (in this case, temperature values) from the dataset.

3. 1:

- The number of features in each time step. Here, each time step contains 1 feature, which is the temperature.

Input 16:

```
import matplotlib.pyplot as plt

import seaborn as sns

from tensorflow.keras.layers import Dense, RepeatVector, LSTM, Dropout

from tensorflow.keras.layers import Flatten, Conv1D, MaxPooling1D

from tensorflow.keras.layers import Bidirectional, Dropout

from tensorflow.keras.models import Sequential

from tensorflow.keras.utils import plot_model
```

Code explanation:

This code prepares the necessary tools for building and visualizing a deep learning model using TensorFlow and Keras, particularly for time series forecasting. The libraries and layers imported will be used to create, train, and visualize a neural network model that can handle sequential data, such as weather forecasting.

- matplotlib and seaborn are for data visualization.
- tensorflow.keras.layers imports various layers for building the model.
- Sequential is used to construct the model in a straightforward manner.
- plot_model allows for visualizing the architecture of the neural network.

Input 17:

```
model = Sequential()
```

```
model.add(LSTM(50, return_sequences= True, input_shape= (20,1)))
model.add(LSTM(50, return_sequences=True))
model.add(LSTM(50))
model.add(Dense(1))
model.compile(loss = 'mean_squared_error', optimizer = 'adam')
model = Sequential()
model.add(LSTM(50, return_sequences= True, input_shape= (20,1)))
model.add(LSTM(50, return_sequences=True))
model.add(LSTM(50))
model.add(Dense(1))
model.compile(loss = 'mean_squared_error', optimizer = 'adam')
model.fit(x_train,y_train,epochs=300, verbose=1, )
```

Epoch 1/300

8/8 [=====] - 5s 31ms/step - loss: 0.2052

Epoch 2/300

8/8 [=====] - 0s 29ms/step - loss: 0.0767

Epoch 3/300

8/8 [=====] - 0s 29ms/step - loss: 0.0594

Epoch 4/300

8/8 [=====] - 0s 29ms/step - loss: 0.0536

Epoch 5/300

8/8 [=====] - 0s 28ms/step - loss: 0.0514

Epoch 6/300

8/8 [=====] - 0s 28ms/step - loss: 0.0495

Epoch 7/300

8/8 [=====] - 0s 29ms/step - loss: 0.0498

Epoch 8/300

8/8 [=====] - 0s 30ms/step - loss: 0.0448

Epoch 9/300

8/8 [=====] - 0s 29ms/step - loss: 0.0421

Epoch 10/300

8/8 [=====] - 0s 26ms/step - loss: 0.0406

Epoch 11/300

8/8 [=====] - 0s 29ms/step - loss: 0.0402

Epoch 12/300

8/8 [=====] - 0s 30ms/step - loss: 0.0397

Epoch 13/300

8/8 [=====] - 0s 29ms/step - loss: 0.0310

Epoch 14/300

8/8 [=====] - 0s 29ms/step - loss: 0.0237

Epoch 15/300

8/8 [=====] - 0s 30ms/step - loss: 0.0228

Epoch 16/300

8/8 [=====] - 0s 29ms/step - loss: 0.0221

Epoch 17/300

8/8 [=====] - 0s 29ms/step - loss: 0.0162

Epoch 18/300

8/8 [=====] - 0s 28ms/step - loss: 0.0175

Epoch 19/300

8/8 [=====] - 0s 29ms/step - loss: 0.0152

Epoch 20/300

8/8 [=====] - 0s 29ms/step - loss: 0.0125

Epoch 21/300

8/8 [=====] - 0s 33ms/step - loss: 0.0143

Epoch 22/300

8/8 [=====] - 0s 29ms/step - loss: 0.0165

Epoch 23/300

8/8 [=====] - 0s 32ms/step - loss: 0.0142

Epoch 24/300

8/8 [=====] - 0s 30ms/step - loss: 0.0114

Epoch 25/300

8/8 [=====] - 0s 29ms/step - loss: 0.0102

Epoch 26/300

8/8 [=====] - 0s 29ms/step - loss: 0.0090

Epoch 27/300

8/8 [=====] - 0s 28ms/step - loss: 0.0098

Epoch 28/300

8/8 [=====] - 0s 29ms/step - loss: 0.0094

Epoch 29/300

8/8 [=====] - 0s 34ms/step - loss: 0.0092

Epoch 30/300

8/8 [=====] - 0s 28ms/step - loss: 0.0091

Epoch 31/300

8/8 [=====] - 0s 30ms/step - loss: 0.0094

Epoch 32/300

8/8 [=====] - 0s 29ms/step - loss: 0.0081

Epoch 33/300

8/8 [=====] - 0s 28ms/step - loss: 0.0095

Epoch 34/300

8/8 [=====] - 0s 28ms/step - loss: 0.0072

Epoch 35/300

8/8 [=====] - 0s 28ms/step - loss: 0.0082

Epoch 36/300

8/8 [=====] - 0s 29ms/step - loss: 0.0079

Epoch 37/300

8/8 [=====] - 0s 28ms/step - loss: 0.0083

Epoch 38/300

8/8 [=====] - 0s 30ms/step - loss: 0.0084

Epoch 39/300

8/8 [=====] - 0s 32ms/step - loss: 0.0109

Epoch 40/300

8/8 [=====] - 0s 28ms/step - loss: 0.0083

Epoch 41/300

8/8 [=====] - 0s 28ms/step - loss: 0.0076

Epoch 42/300

8/8 [=====] - 0s 29ms/step - loss: 0.0091

Epoch 43/300

8/8 [=====] - 0s 27ms/step - loss: 0.0074

Epoch 44/300

8/8 [=====] - 0s 29ms/step - loss: 0.0090

Epoch 45/300

8/8 [=====] - 0s 28ms/step - loss: 0.0075

Epoch 46/300

8/8 [=====] - 0s 30ms/step - loss: 0.0070

Epoch 47/300

8/8 [=====] - 0s 30ms/step - loss: 0.0084

Epoch 48/300

8/8 [=====] - 0s 33ms/step - loss: 0.0067

Epoch 49/300

8/8 [=====] - 0s 37ms/step - loss: 0.0099

Epoch 50/300

8/8 [=====] - 0s 36ms/step - loss: 0.0071

Epoch 51/300

8/8 [=====] - 0s 35ms/step - loss: 0.0079

Epoch 52/300

8/8 [=====] - 0s 33ms/step - loss: 0.0080

Epoch 53/300

8/8 [=====] - 0s 36ms/step - loss: 0.0066

Epoch 54/300

8/8 [=====] - 0s 36ms/step - loss: 0.0071

Epoch 55/300

8/8 [=====] - 0s 29ms/step - loss: 0.0061

Epoch 56/300

8/8 [=====] - 0s 29ms/step - loss: 0.0065

Epoch 57/300

8/8 [=====] - 0s 30ms/step - loss: 0.0067

Epoch 58/300

8/8 [=====] - 0s 31ms/step - loss: 0.0079

Epoch 59/300

8/8 [=====] - 0s 35ms/step - loss: 0.0067

Epoch 60/300

8/8 [=====] - 0s 28ms/step - loss: 0.0074

Epoch 61/300

8/8 [=====] - 0s 30ms/step - loss: 0.0066

Epoch 62/300

8/8 [=====] - 0s 30ms/step - loss: 0.0060

Epoch 63/300

8/8 [=====] - 0s 27ms/step - loss: 0.0055

Epoch 64/300

8/8 [=====] - 0s 28ms/step - loss: 0.0061

Epoch 65/300

8/8 [=====] - 0s 28ms/step - loss: 0.0062

Epoch 66/300

8/8 [=====] - 0s 29ms/step - loss: 0.0067

Epoch 67/300

8/8 [=====] - 0s 29ms/step - loss: 0.0056

Epoch 68/300

8/8 [=====] - 0s 29ms/step - loss: 0.0068

Epoch 69/300

8/8 [=====] - 0s 29ms/step - loss: 0.0050

Epoch 70/300

8/8 [=====] - 0s 29ms/step - loss: 0.0066

Epoch 71/300

8/8 [=====] - 0s 30ms/step - loss: 0.0045

Epoch 72/300

8/8 [=====] - 0s 30ms/step - loss: 0.0057

Epoch 73/300

8/8 [=====] - 0s 39ms/step - loss: 0.0056

Epoch 74/300

8/8 [=====] - 0s 30ms/step - loss: 0.0067

Epoch 75/300

8/8 [=====] - 0s 30ms/step - loss: 0.0070

Epoch 76/300

8/8 [=====] - 0s 29ms/step - loss: 0.0069

Epoch 77/300

8/8 [=====] - 0s 30ms/step - loss: 0.0079

Epoch 78/300

8/8 [=====] - 0s 28ms/step - loss: 0.0061

Epoch 79/300

8/8 [=====] - 0s 28ms/step - loss: 0.0075

Epoch 80/300

8/8 [=====] - 0s 28ms/step - loss: 0.0059

Epoch 81/300

8/8 [=====] - 0s 29ms/step - loss: 0.0055

Epoch 82/300

8/8 [=====] - 0s 29ms/step - loss: 0.0067

Epoch 83/300

8/8 [=====] - 0s 28ms/step - loss: 0.0066

Epoch 84/300

8/8 [=====] - 0s 28ms/step - loss: 0.0060

Epoch 85/300

8/8 [=====] - 0s 28ms/step - loss: 0.0053

Epoch 86/300

8/8 [=====] - 0s 29ms/step - loss: 0.0051

Epoch 87/300

8/8 [=====] - 0s 29ms/step - loss: 0.0048

Epoch 88/300

8/8 [=====] - 0s 29ms/step - loss: 0.0049

Epoch 89/300

8/8 [=====] - 0s 27ms/step - loss: 0.0048

Epoch 90/300

8/8 [=====] - 0s 30ms/step - loss: 0.0048

Epoch 91/300

8/8 [=====] - 0s 30ms/step - loss: 0.0045

Epoch 92/300

8/8 [=====] - 0s 28ms/step - loss: 0.0060

Epoch 93/300

8/8 [=====] - 0s 30ms/step - loss: 0.0049

Epoch 94/300

8/8 [=====] - 0s 29ms/step - loss: 0.0047

Epoch 95/300

8/8 [=====] - 0s 29ms/step - loss: 0.0041

Epoch 96/300

8/8 [=====] - 0s 33ms/step - loss: 0.0040

Epoch 97/300

8/8 [=====] - 0s 34ms/step - loss: 0.0049

Epoch 98/300

8/8 [=====] - 0s 32ms/step - loss: 0.0041

Epoch 99/300

8/8 [=====] - 0s 37ms/step - loss: 0.0042

Epoch 100/300

8/8 [=====] - 0s 34ms/step - loss: 0.0038

Epoch 101/300

8/8 [=====] - 0s 34ms/step - loss: 0.0040

Epoch 102/300

8/8 [=====] - 0s 39ms/step - loss: 0.0040

Epoch 103/300

8/8 [=====] - 0s 31ms/step - loss: 0.0046

Epoch 104/300

8/8 [=====] - 0s 36ms/step - loss: 0.0041

Epoch 105/300

8/8 [=====] - 0s 38ms/step - loss: 0.0043

Epoch 106/300

8/8 [=====] - 0s 30ms/step - loss: 0.0048

Epoch 107/300

8/8 [=====] - 0s 31ms/step - loss: 0.0036

Epoch 108/300

8/8 [=====] - 0s 32ms/step - loss: 0.0040

Epoch 109/300

8/8 [=====] - 0s 35ms/step - loss: 0.0040

Epoch 110/300

8/8 [=====] - 0s 29ms/step - loss: 0.0038

Epoch 111/300

8/8 [=====] - 0s 28ms/step - loss: 0.0029

Epoch 112/300

8/8 [=====] - 0s 31ms/step - loss: 0.0037

Epoch 113/300

8/8 [=====] - 0s 31ms/step - loss: 0.0035

Epoch 114/300

8/8 [=====] - 0s 30ms/step - loss: 0.0029

Epoch 115/300

8/8 [=====] - 0s 33ms/step - loss: 0.0030

Epoch 116/300

8/8 [=====] - 0s 33ms/step - loss: 0.0037

Epoch 117/300

8/8 [=====] - 0s 37ms/step - loss: 0.0040

Epoch 118/300

8/8 [=====] - 0s 33ms/step - loss: 0.0032

Epoch 119/300

8/8 [=====] - 0s 36ms/step - loss: 0.0035

Epoch 120/300

8/8 [=====] - 0s 29ms/step - loss: 0.0032

Epoch 121/300

8/8 [=====] - 0s 32ms/step - loss: 0.0035

Epoch 122/300

8/8 [=====] - 0s 33ms/step - loss: 0.0036

Epoch 123/300

8/8 [=====] - 0s 34ms/step - loss: 0.0037

Epoch 124/300

8/8 [=====] - 0s 30ms/step - loss: 0.0040

Epoch 125/300

8/8 [=====] - 0s 31ms/step - loss: 0.0034

Epoch 126/300

8/8 [=====] - 0s 33ms/step - loss: 0.0046

Epoch 127/300

8/8 [=====] - 0s 31ms/step - loss: 0.0045

Epoch 128/300

8/8 [=====] - 0s 34ms/step - loss: 0.0041

Epoch 129/300

8/8 [=====] - 0s 31ms/step - loss: 0.0034

Epoch 130/300

8/8 [=====] - 0s 34ms/step - loss: 0.0029

Epoch 131/300

8/8 [=====] - 0s 33ms/step - loss: 0.0043

Epoch 132/300

8/8 [=====] - 0s 31ms/step - loss: 0.0038

Epoch 133/300

8/8 [=====] - 0s 30ms/step - loss: 0.0040

Epoch 134/300

8/8 [=====] - 0s 28ms/step - loss: 0.0035

Epoch 135/300

8/8 [=====] - 0s 37ms/step - loss: 0.0036

Epoch 136/300

8/8 [=====] - 0s 29ms/step - loss: 0.0040

Epoch 137/300

8/8 [=====] - 0s 29ms/step - loss: 0.0038

Epoch 138/300

8/8 [=====] - 0s 29ms/step - loss: 0.0035

Epoch 139/300

8/8 [=====] - 0s 30ms/step - loss: 0.0043

Epoch 140/300

8/8 [=====] - 0s 30ms/step - loss: 0.0038

Epoch 141/300

8/8 [=====] - 0s 30ms/step - loss: 0.0032

Epoch 142/300

8/8 [=====] - 0s 29ms/step - loss: 0.0033

Epoch 143/300

8/8 [=====] - 0s 32ms/step - loss: 0.0026

Epoch 144/300

8/8 [=====] - 0s 32ms/step - loss: 0.0031

Epoch 145/300

8/8 [=====] - 0s 32ms/step - loss: 0.0030

Epoch 146/300

8/8 [=====] - 0s 30ms/step - loss: 0.0032

Epoch 147/300

8/8 [=====] - 0s 31ms/step - loss: 0.0037

Epoch 148/300

8/8 [=====] - 0s 38ms/step - loss: 0.0032

Epoch 149/300

8/8 [=====] - 0s 31ms/step - loss: 0.0026

Epoch 150/300

8/8 [=====] - 0s 31ms/step - loss: 0.0029

Epoch 151/300

8/8 [=====] - 0s 30ms/step - loss: 0.0032

Epoch 152/300

8/8 [=====] - 0s 30ms/step - loss: 0.0029

Epoch 153/300

8/8 [=====] - 0s 30ms/step - loss: 0.0027

Epoch 154/300

8/8 [=====] - 0s 29ms/step - loss: 0.0032

Epoch 155/300

8/8 [=====] - 0s 28ms/step - loss: 0.0029

Epoch 156/300

8/8 [=====] - 0s 30ms/step - loss: 0.0032

Epoch 157/300

8/8 [=====] - 0s 28ms/step - loss: 0.0027

Epoch 158/300

8/8 [=====] - 0s 29ms/step - loss: 0.0034

Epoch 159/300

8/8 [=====] - 0s 30ms/step - loss: 0.0024

Epoch 160/300

8/8 [=====] - 0s 33ms/step - loss: 0.0035

Epoch 161/300

8/8 [=====] - 0s 31ms/step - loss: 0.0033

Epoch 162/300

8/8 [=====] - 0s 30ms/step - loss: 0.0034

Epoch 163/300

8/8 [=====] - 0s 29ms/step - loss: 0.0032

Epoch 164/300

8/8 [=====] - 0s 29ms/step - loss: 0.0032

Epoch 165/300

8/8 [=====] - 0s 29ms/step - loss: 0.0029

Epoch 166/300

8/8 [=====] - 0s 28ms/step - loss: 0.0028

Epoch 167/300

8/8 [=====] - 0s 29ms/step - loss: 0.0029

Epoch 168/300

8/8 [=====] - 0s 30ms/step - loss: 0.0028

Epoch 169/300

8/8 [=====] - 0s 28ms/step - loss: 0.0033

Epoch 170/300

8/8 [=====] - 0s 29ms/step - loss: 0.0035

Epoch 171/300

8/8 [=====] - 0s 29ms/step - loss: 0.0029

Epoch 172/300

8/8 [=====] - 0s 29ms/step - loss: 0.0027

Epoch 173/300

8/8 [=====] - 0s 30ms/step - loss: 0.0029

Epoch 174/300

8/8 [=====] - 0s 30ms/step - loss: 0.0031

Epoch 175/300

8/8 [=====] - 0s 29ms/step - loss: 0.0037

Epoch 176/300

8/8 [=====] - 0s 29ms/step - loss: 0.0032

Epoch 177/300

8/8 [=====] - 0s 29ms/step - loss: 0.0034

Epoch 178/300

8/8 [=====] - 0s 27ms/step - loss: 0.0039

Epoch 179/300

8/8 [=====] - 0s 29ms/step - loss: 0.0035

Epoch 180/300

8/8 [=====] - 0s 31ms/step - loss: 0.0030

Epoch 181/300

8/8 [=====] - 0s 29ms/step - loss: 0.0036

Epoch 182/300

8/8 [=====] - 0s 29ms/step - loss: 0.0029

Epoch 183/300

8/8 [=====] - 0s 36ms/step - loss: 0.0029

Epoch 184/300

8/8 [=====] - 0s 34ms/step - loss: 0.0026

Epoch 185/300

8/8 [=====] - 0s 31ms/step - loss: 0.0025

Epoch 186/300

8/8 [=====] - 0s 31ms/step - loss: 0.0026

Epoch 187/300

8/8 [=====] - 0s 32ms/step - loss: 0.0025

Epoch 188/300

8/8 [=====] - 0s 29ms/step - loss: 0.0039

Epoch 189/300

8/8 [=====] - 0s 28ms/step - loss: 0.0031

Epoch 190/300

8/8 [=====] - 0s 29ms/step - loss: 0.0031

Epoch 191/300

8/8 [=====] - 0s 28ms/step - loss: 0.0028

Epoch 192/300

8/8 [=====] - 0s 27ms/step - loss: 0.0035

Epoch 193/300

8/8 [=====] - 0s 28ms/step - loss: 0.0033

Epoch 194/300

8/8 [=====] - 0s 28ms/step - loss: 0.0024

Epoch 195/300

8/8 [=====] - 0s 29ms/step - loss: 0.0034

Epoch 196/300

8/8 [=====] - 0s 29ms/step - loss: 0.0034

Epoch 197/300

8/8 [=====] - 0s 29ms/step - loss: 0.0027

Epoch 198/300

8/8 [=====] - 0s 29ms/step - loss: 0.0029

Epoch 199/300

8/8 [=====] - 0s 28ms/step - loss: 0.0028

Epoch 200/300

8/8 [=====] - 0s 29ms/step - loss: 0.0030

Epoch 201/300

8/8 [=====] - 0s 28ms/step - loss: 0.0033

Epoch 202/300

8/8 [=====] - 0s 29ms/step - loss: 0.0031

Epoch 203/300

8/8 [=====] - 0s 28ms/step - loss: 0.0032

Epoch 204/300

8/8 [=====] - 0s 29ms/step - loss: 0.0028

Epoch 205/300

8/8 [=====] - 0s 29ms/step - loss: 0.0030

Epoch 206/300

8/8 [=====] - 0s 33ms/step - loss: 0.0030

Epoch 207/300

8/8 [=====] - 0s 33ms/step - loss: 0.0027

Epoch 208/300

8/8 [=====] - 0s 28ms/step - loss: 0.0026

Epoch 209/300

8/8 [=====] - 0s 28ms/step - loss: 0.0029

Epoch 210/300

8/8 [=====] - 0s 28ms/step - loss: 0.0026

Epoch 211/300

8/8 [=====] - 0s 28ms/step - loss: 0.0027

Epoch 212/300

8/8 [=====] - 0s 29ms/step - loss: 0.0032

Epoch 213/300

8/8 [=====] - 0s 27ms/step - loss: 0.0037

Epoch 214/300

8/8 [=====] - 0s 28ms/step - loss: 0.0030

Epoch 215/300

8/8 [=====] - 0s 30ms/step - loss: 0.0042

Epoch 216/300

8/8 [=====] - 0s 29ms/step - loss: 0.0033

Epoch 217/300

8/8 [=====] - 0s 30ms/step - loss: 0.0035

Epoch 218/300

8/8 [=====] - 0s 30ms/step - loss: 0.0037

Epoch 219/300

8/8 [=====] - 0s 29ms/step - loss: 0.0030

Epoch 220/300

8/8 [=====] - 0s 28ms/step - loss: 0.0040

Epoch 221/300

8/8 [=====] - 0s 27ms/step - loss: 0.0030

Epoch 222/300

8/8 [=====] - 0s 28ms/step - loss: 0.0028

Epoch 223/300

8/8 [=====] - 0s 29ms/step - loss: 0.0025

Epoch 224/300

8/8 [=====] - 0s 29ms/step - loss: 0.0032

Epoch 225/300

8/8 [=====] - 0s 30ms/step - loss: 0.0035

Epoch 226/300

8/8 [=====] - 0s 28ms/step - loss: 0.0037

Epoch 227/300

8/8 [=====] - 0s 28ms/step - loss: 0.0034

Epoch 228/300

8/8 [=====] - 0s 31ms/step - loss: 0.0027

Epoch 229/300

8/8 [=====] - 0s 32ms/step - loss: 0.0029

Epoch 230/300

8/8 [=====] - 0s 31ms/step - loss: 0.0035

Epoch 231/300

8/8 [=====] - 0s 32ms/step - loss: 0.0034

Epoch 232/300

8/8 [=====] - 0s 32ms/step - loss: 0.0029

Epoch 233/300

8/8 [=====] - 0s 31ms/step - loss: 0.0026

Epoch 234/300

8/8 [=====] - 0s 29ms/step - loss: 0.0035

Epoch 235/300

8/8 [=====] - 0s 30ms/step - loss: 0.0029

Epoch 236/300

8/8 [=====] - 0s 29ms/step - loss: 0.0030

Epoch 237/300

8/8 [=====] - 0s 29ms/step - loss: 0.0029

Epoch 238/300

8/8 [=====] - 0s 29ms/step - loss: 0.0028

Epoch 239/300

8/8 [=====] - 0s 29ms/step - loss: 0.0030

Epoch 240/300

8/8 [=====] - 0s 29ms/step - loss: 0.0030

Epoch 241/300

8/8 [=====] - 0s 28ms/step - loss: 0.0028

Epoch 242/300

8/8 [=====] - 0s 30ms/step - loss: 0.0026

Epoch 243/300

8/8 [=====] - 0s 31ms/step - loss: 0.0033

Epoch 244/300

8/8 [=====] - 0s 29ms/step - loss: 0.0033

Epoch 245/300

8/8 [=====] - 0s 30ms/step - loss: 0.0039

Epoch 246/300

8/8 [=====] - 0s 30ms/step - loss: 0.0026

Epoch 247/300

8/8 [=====] - 0s 28ms/step - loss: 0.0030

Epoch 248/300

8/8 [=====] - 0s 30ms/step - loss: 0.0032

Epoch 249/300

8/8 [=====] - 0s 28ms/step - loss: 0.0024

Epoch 250/300

8/8 [=====] - 0s 32ms/step - loss: 0.0035

Epoch 251/300

8/8 [=====] - 0s 31ms/step - loss: 0.0032

Epoch 252/300

8/8 [=====] - 0s 31ms/step - loss: 0.0036

Epoch 253/300

8/8 [=====] - 0s 37ms/step - loss: 0.0037

Epoch 254/300

8/8 [=====] - 0s 29ms/step - loss: 0.0032

Epoch 255/300

8/8 [=====] - 0s 29ms/step - loss: 0.0035

Epoch 256/300

8/8 [=====] - 0s 28ms/step - loss: 0.0034

Epoch 257/300

8/8 [=====] - 0s 29ms/step - loss: 0.0027

Epoch 258/300

8/8 [=====] - 0s 32ms/step - loss: 0.0027

Epoch 259/300

8/8 [=====] - 0s 32ms/step - loss: 0.0026

Epoch 260/300

8/8 [=====] - 0s 28ms/step - loss: 0.0026

Epoch 261/300

8/8 [=====] - 0s 30ms/step - loss: 0.0030

Epoch 262/300

8/8 [=====] - 0s 29ms/step - loss: 0.0025

Epoch 263/300

8/8 [=====] - 0s 28ms/step - loss: 0.0024

Epoch 264/300

8/8 [=====] - 0s 28ms/step - loss: 0.0029

Epoch 265/300

8/8 [=====] - 0s 28ms/step - loss: 0.0029

Epoch 266/300

8/8 [=====] - 0s 29ms/step - loss: 0.0027

Epoch 267/300

8/8 [=====] - 0s 28ms/step - loss: 0.0028

Epoch 268/300

8/8 [=====] - 0s 30ms/step - loss: 0.0025

Epoch 269/300

8/8 [=====] - 0s 28ms/step - loss: 0.0032

Epoch 270/300

8/8 [=====] - 0s 29ms/step - loss: 0.0030

Epoch 271/300

8/8 [=====] - 0s 29ms/step - loss: 0.0032

Epoch 272/300

8/8 [=====] - 0s 29ms/step - loss: 0.0028

Epoch 273/300

8/8 [=====] - 0s 30ms/step - loss: 0.0030

Epoch 274/300

8/8 [=====] - 0s 28ms/step - loss: 0.0034

Epoch 275/300

8/8 [=====] - 0s 28ms/step - loss: 0.0025

Epoch 276/300

8/8 [=====] - 0s 29ms/step - loss: 0.0025

Epoch 277/300

8/8 [=====] - 0s 29ms/step - loss: 0.0034

Epoch 278/300

8/8 [=====] - 0s 29ms/step - loss: 0.0029

Epoch 279/300

8/8 [=====] - 0s 29ms/step - loss: 0.0026

Epoch 280/300

8/8 [=====] - 0s 28ms/step - loss: 0.0029

Epoch 281/300

8/8 [=====] - 0s 27ms/step - loss: 0.0023

Epoch 282/300

8/8 [=====] - 0s 29ms/step - loss: 0.0033

Epoch 283/300

8/8 [=====] - 0s 30ms/step - loss: 0.0024

Epoch 284/300

8/8 [=====] - 0s 28ms/step - loss: 0.0030

Epoch 285/300

8/8 [=====] - 0s 29ms/step - loss: 0.0025

Epoch 286/300

8/8 [=====] - 0s 29ms/step - loss: 0.0025

Epoch 287/300

8/8 [=====] - 0s 28ms/step - loss: 0.0022

Epoch 288/300

8/8 [=====] - 0s 27ms/step - loss: 0.0027

Epoch 289/300

8/8 [=====] - 0s 30ms/step - loss: 0.0030

Epoch 290/300

8/8 [=====] - 0s 27ms/step - loss: 0.0039

Epoch 291/300

8/8 [=====] - 0s 29ms/step - loss: 0.0040

Epoch 292/300

8/8 [=====] - 0s 29ms/step - loss: 0.0033

Epoch 293/300

8/8 [=====] - 0s 29ms/step - loss: 0.0033

Epoch 294/300

8/8 [=====] - 0s 31ms/step - loss: 0.0032

Epoch 295/300

8/8 [=====] - 0s 34ms/step - loss: 0.0030

Epoch 296/300

8/8 [=====] - 0s 28ms/step - loss: 0.0041

Epoch 297/300

8/8 [=====] - 0s 28ms/step - loss: 0.0031

Epoch 298/300

8/8 [=====] - 0s 29ms/step - loss: 0.0029

Epoch 299/300

8/8 [=====] - 0s 35ms/step - loss: 0.0031

Epoch 300/300

8/8 [=====] - 0s 29ms/step - loss: 0.0033

Output:

<tensorflow.python.keras.callbacks.History at 0x7f347005e350>

Code explanation:

□ `model = Sequential():`

- Initializes a sequential model. The sequential model is a linear stack of layers, which is suitable for most deep learning tasks.

□ `model.add(LSTM(50, return_sequences= True, input_shape= (20,1))):`

- Adds the first LSTM layer with 50 units.

- `return_sequences=True` ensures that the LSTM layer returns the full sequence output instead of just the last output, which is necessary when stacking multiple LSTM layers.
 - `input_shape=(20,1)` specifies that each input sequence has 20 time steps, and each time step has 1 feature (e.g., temperature).
- `model.add(LSTM(50, return_sequences=True))`:
- Adds a second LSTM layer with 50 units, also returning sequences.
- `model.add(LSTM(50))`:
- Adds a third LSTM layer with 50 units, but this time without returning sequences (returning only the final output).
- `model.add(Dense(1))`:
- Adds a Dense layer with 1 unit, which will be the output layer. This layer produces the final prediction (e.g., the next temperature value).
- `model.compile(loss = 'mean_squared_error', optimizer = 'adam')`:
- Compiles the model by defining the loss function and optimizer.
 - `loss='mean_squared_error'` indicates that the model will use Mean Squared Error (MSE) as the loss function, which is common for regression tasks.
 - `optimizer='adam'` specifies the Adam optimizer, which is efficient and widely used in training neural networks.
- `model.fit(x_train, y_train, epochs=300, verbose=1)`:
- Trains the model using the training data (`x_train`, `y_train`).
 - `epochs=300` means the model will go through the entire dataset 300 times.
 - `verbose=1` will display a progress bar showing the training process.

Output explanation:

- The output `<tensorflow.python.keras.callbacks.History at 0x7f347005e350>` simply indicates that the model has been trained and that a History object has been created to store the training metrics.
- You can use this object to analyze the training process, such as checking how the loss decreased over time.

Input 18:

```
# model.evaluate(x_test, y_test)
```

Output:

Predictions and model evaluation

Input 19:

```
print("Predicted Value",model.predict(x_train)[4][0])
```

```
print("Expected value",y_train[4][0])
```

Output:

Predicted Value 0.7321534

Expected value 0.7083333333333334

Code explanation:

- `model.predict(x_train)`:
 - This line uses the trained LSTM model to make predictions on the training data (`x_train`).
 - The output of `model.predict(x_train)` is an array of predicted values for all samples in the `x_train` dataset.
- `model.predict(x_train)[4][0]`:
 - This retrieves the predicted value for the 5th sample (index 4) in the `x_train` dataset.

- `[4][0]` accesses the first element of the 5th prediction.

□ `y_train[4][0]`:

- This retrieves the actual expected value (target) for the 5th sample in the `y_train` dataset.
- Since `y_train` is also an array, `[4][0]` accesses the first element of the 5th expected value.

□ `print()`:

- The `print()` statements output the predicted and expected values to the console, allowing you to see how close the model's prediction is to the actual value.

Output explanation:

□ Predicted Value 0.7321534:

- This is the value predicted by the LSTM model for the 5th sample in the `x_train` dataset after training.
- The model outputs a normalized value (since the data was scaled using `MinMaxScaler` earlier).

□ Expected Value 0.7083333333333334:

- This is the actual value for the 5th sample in the `y_train` dataset. It represents the normalized expected value for the corresponding input sequence in `x_train`.

Input 20:

```
predictions=model.predict(x_test)
```

```
print("Predicted Value",predictions[2][0])
```

```
print("Expected Value",y_test[2][0])
```

Output:

```
Predicted Value 0.42669037
```

```
Expected Value 0.41666666666666663
```

Code explanation:

□ `predictions = model.predict(x_test):`

- This line uses the trained model to predict values based on the `x_test` data.
- `x_test` is a portion of the data that was not used during training, allowing you to test how well the model generalizes to unseen data.
- The output, `predictions`, is an array where each element corresponds to a predicted value for each instance in `x_test`.

□ `print("Predicted Value", predictions[2][0]):`

- This line prints the predicted value for the third instance in the test set (index 2).
- `predictions[2][0]` accesses the first element of the third prediction. Since this is a time-series prediction, the result is usually a single value.

□ `print("Expected Value", y_test[2][0]):`

- This line prints the actual (expected) value for the third instance in the test set.
- `y_test[2][0]` accesses the first element of the third actual value in the test set.

Output explanation:

□ Predicted Value: 0.42669037

- This value represents the result produced by the trained model for the third instance in the `x_test` dataset.
- The value 0.42669037 indicates the model's best estimate for the target variable based on the input features it received.

□ Expected Value: 0.41666666666666663

- This value represents the actual or true value from the `y_test` dataset, which corresponds to the same instance that was used for the prediction.
- The expected value 0.41666666666666663 is the ground truth that the model was trying to predict.

Input 21:

```
predictions.shape
```

Output:

```
(105, 1)
```

Code explanation:

This code is used to get the shape (or dimensions) of the predictions array.

Output explanation:

The shape (105, 1) confirms that the model generated 105 predictions, each with a single output value, which aligns with the expectation for a univariate time series prediction task.

Input 22:

```
y_test.shape
```

Output:

```
(105, 1)
```

Code explanation:

This code retrieves the shape (dimensions) of the y_test array.

Output explanation:

The output (105, 1) shows that y_test contains 105 actual values, each with a single output. This matches the shape of the predictions array, allowing for a direct comparison between predicted and actual values during model evaluation.

Input 23:

```
pred_df=pd.DataFrame(predictions)
```

```
pred_df['TrueValues']=y_test
```

```
pred_df_new = pred_df.rename(columns={ 0: 'Predictions'})
```

```
pred_df_new
```

Output:

Predictions	True Values	
0	0.469617	0.541667
1	0.574576	0.416667
2	0.426690	0.416667
3	0.434639	0.458333
4	0.505478	0.500000
...
100	0.241969	0.291667
101	0.293324	0.208333
102	0.245331	0.250000
103	0.283295	0.250000
104	0.311354	0.291667

105 rows × 2 columns

Code explanation:

- Convert Predictions to DataFrame (`pred_df = pd.DataFrame(predictions)`): The predictions array is turned into a DataFrame, with each prediction as a row.
- Add Actual Values (`pred_df['TrueValues'] = y_test`): A new column, TrueValues, is added to this DataFrame, containing the actual values (`y_test`).
- Rename Columns (`pred_df_new = pred_df.rename(columns={0: 'Predictions'})`): The predictions column is renamed from 0 to Predictions for clarity.
- Result (`pred_df_new`): The final DataFrame (`pred_df_new`) displays the predicted values alongside the true values, making it easier to evaluate the model's accuracy.

Output explanation:

The output of the code, presented as a DataFrame, shows a comparison between the predicted values generated by the model and the actual true values. Here's an explanation:

Output Structure:

- Predictions: This column contains the predicted temperature values normalized between 0 and 1, generated by the LSTM model.
- TrueValues: This column holds the corresponding actual temperature values (also normalized) that the model was trying to predict.

Interpretation:

- The DataFrame has 105 rows, each representing a time step where a prediction was made.
- For instance:
 - The first row shows a predicted value of 0.469617 compared to the true value of 0.541667.
 - The second row shows a predicted value of 0.574576 compared to the true value of 0.416667.
- The difference between the values in each row indicates the model's accuracy. If the predictions are close to the true values, the model is performing well.

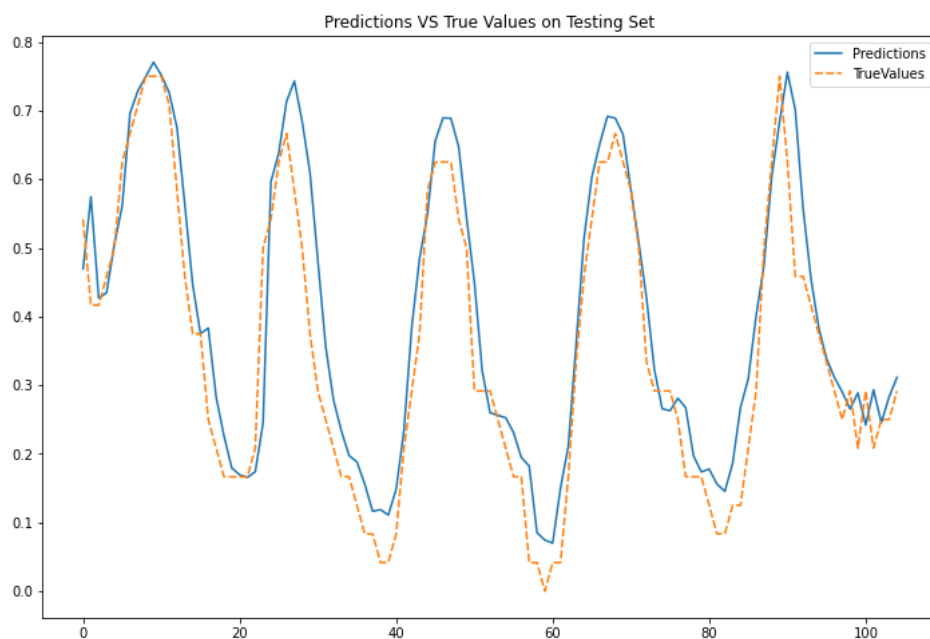
This table allows you to see how well the model's predictions match the actual data across all the time steps in the test set.

Input 24:

```
plt.figure(figsize=(12,8))  
  
sns.lineplot(data= pred_df_new)  
  
plt.title("Predictions VS True Values on Testing Set")
```

Output:

```
Text(0.5, 1.0, 'Predictions VS True Values on Testing Set')
```



Code explanation:

- `plt.figure(figsize=(12,8)):`
 - Creates a new figure for the plot with a size of 12 inches by 8 inches.
- `sns.lineplot(data=pred_df_new):`
 - Uses Seaborn's lineplot function to plot the data from `pred_df_new`.

- The DataFrame `pred_df_new` contains two columns: "Predictions" and "TrueValues".
- Seaborn will plot both columns on the same graph, making it easy to compare the two series.

□ `plt.title("Predictions VS True Values on Testing Set"):`

- Adds a title to the plot, indicating that the graph shows a comparison between the model's predictions and the true values.

Output explanation:

□ Line Plot:

- Blue Line: Represents the predicted values generated by the LSTM model.
- Orange Dotted Line: Represents the true values from the test dataset.

□ Observations:

- The blue and orange lines are closely aligned for most of the data points, indicating that the model predictions closely match the true values.
- Some deviations occur, where the predictions slightly overestimate or underestimate the true values, which is expected in most real-world models.

□ Title:

- The title of the plot is "Predictions VS True Values on Testing Set", clearly indicating what the plot is showing. The text object `Text(0.5, 1.0, 'Predictions VS True Values on Testing Set')` is related to setting this title programmatically.

The plot visually confirms that the model has learned the patterns in the data well, as the predictions (blue line) follow the true values (orange dotted line) closely. The proximity of the lines suggests that the model performs well on the test set, with relatively small prediction errors.

Input 25:

`data_scaled.shape`

Output:

(362, 1)

Code explanation:

□ `data_scaled`: This is a variable that typically represents a NumPy array, pandas DataFrame, or similar data structure containing scaled data (often used in machine learning to normalize or standardize features).

□ `.shape`: This attribute is used to get the dimensions of the array or DataFrame. For a NumPy array, it returns a tuple representing the size of each dimension.

Output explanation:

Output (362, 1):

- 362: This is the number of rows in the `data_scaled` array or Data Frame. It indicates that there are 362 individual data points or samples.
- 1: This is the number of columns in the `data_scaled` array or Data Frame. It shows that there is only one feature or variable for each data point.

Input 26:

```
x_input=data_scaled[:20]
```

```
x_input.shape
```

Output:

(20, 1)

Code explanation:

□ `data_scaled[:20]`:

- This is a slicing operation. It extracts a subset of the `data_scaled` array or DataFrame.

- `:20` means you are selecting all rows from the beginning up to (but not including) the 20th row.
- This operation results in a new array or DataFrame containing the first 20 rows of `data_scaled`.

□ `x_input`:

- This variable will now hold the subset of `data_scaled` consisting of the first 20 rows.

□ `x_input.shape`:

- This returns the dimensions of the `x_input` array or DataFrame.

Output explanation:

- Output (20, 1):
 - 20: This indicates the number of rows in `x_input`. It means `x_input` contains 20 rows, which are the first 20 rows from `data_scaled`.
 - 1: This indicates the number of columns in `x_input`. It shows that there is only one column in this subset, just like in the original `data_scaled`.

So, the shape (20, 1) means that `x_input` is a subset of `data_scaled` containing 20 rows and 1 column.

Input 27:

```
# now reshaping the data
```

```
x_input = x_input.reshape(1, -1)
```

```
x_input.shape
```

Output:

```
(1, 20)
```

Code explanation: `x_input`: This is the variable holding the data that you want to reshape. Before reshaping, it had the shape (20, 1).

□ `.reshape()`: This method is used to change the shape of an array without altering its data. It allows you to specify a new shape for the array.

□ `1`: This specifies that the reshaped array should have 1 row.

□ `-1`: This is a placeholder that means "infer this dimension based on the other dimensions." The `-1` allows NumPy to automatically calculate the number of columns required to fit all the elements of the array based on the specified number of rows.

Output explanation:

- - `1`: This indicates that `x_input` now has 1 row.
 - `20`: This indicates that `x_input` has 20 columns. The reshape operation inferred that 20 columns are needed to accommodate all the elements from the original `x_input`, which had 20 elements (from the 20 rows and 1 column of the original data).

So, after reshaping, `x_input` is transformed from a shape of (20, 1) (20 rows, 1 column) into (1, 20) (1 row, 20 columns). This transformation is often done in machine learning tasks where a specific input shape is required, such as feeding data into a model that expects a single row of features.

Input 28:

```
# x_input = x_input.reshape((1, 20,1))  
  
# yhat = model.predict(x_input, verbose=0)
```

Code explanation:

□ `x_input = x_input.reshape((1, 20, 1))`:

- `x_input`: Initially, this variable was reshaped to (1, 20), which represents a single row with 20 columns.
 - `.reshape((1, 20, 1))`: This reshapes `x_input` to have three dimensions:
 - 1 row (the first dimension)
 - 20 columns (the second dimension)
 - 1 feature (the third dimension)
 - This transformation is often required when working with models that expect inputs in a 3D format, such as LSTM (Long Short-Term Memory) networks used for time series or sequential data. In this context, the three dimensions represent:
 - The number of samples or sequences (1 in this case)
 - The number of time steps or sequence length (20 here)
 - The number of features per time step (1 feature)
- `yhat = model.predict(x_input, verbose=0)`:
- `model`: This refers to a trained machine learning model, likely an LSTM or similar model that expects inputs in the shape specified.
 - `.predict(x_input, verbose=0)`: This method is used to make predictions based on the input data. The arguments are:
 - `x_input`: The reshaped data you are passing into the model.
 - `verbose=0`: This argument controls the verbosity of the output. 0 means no additional output is shown during prediction.

Input 29:

```
temp_input = list(x_input)
```

```
temp_input = temp_input[0].tolist()
```

Output:

creating a function which returns the next 30 days prediction

Code explanation:

□ `temp_input = list(x_input):`

- `x_input`: This is a NumPy array with the shape (1, 20, 1), meaning it has 1 sample, 20 time steps, and 1 feature.
- `list(x_input)`: This converts the `x_input` NumPy array into a list. The resulting list will have the same structure as the array, but as a Python list instead of a NumPy array.

□ `temp_input = temp_input[0].tolist():`

- `temp_input[0]`: Since `temp_input` is a list of arrays (or nested lists), `temp_input[0]` selects the first element of this list. Given the original shape of (1, 20, 1), `temp_input[0]` will be a list or array with shape (20, 1), which contains the time steps and features.
- `.tolist()`: This converts the array or list with shape (20, 1) into a regular Python list. This results in a nested list structure where each of the 20 elements is itself a list containing 1 feature value.

Input 30:

```
# demonstrate prediction for next 10 days
```

```
from numpy import array
```

```
lst_output=[]
```

```
n_steps=20
```

```
i=0
```

```
while(i<30):
```

```

if(len(temp_input)>20):
    #print(temp_input)

    x_input=np.array(temp_input[1:])

    print("{} day input {}".format(i,x_input))

    x_input=x_input.reshape(1,-1)

    x_input = x_input.reshape((1, n_steps, 1))

    #print(x_input)

    yhat = model.predict(x_input, verbose=0)

    print("{} day output {}".format(i,yhat))

    temp_input.extend(yhat[0].tolist())

    temp_input=temp_input[1:]

    #print(temp_input)

    lst_output.extend(yhat.tolist())

    i=i+1
else:

    x_input = x_input.reshape((1, n_steps,1))

    yhat = model.predict(x_input, verbose=0)

    print(yhat[0])

    temp_input.extend(yhat[0].tolist())

#    print(len(temp_input))

    lst_output.extend(yhat.tolist())

    i=i+1

```

```
print(lst_output)
```

Output:

```
[0.8552897]
```

```
1 day input [0.75    0.58333333 0.58333333 0.54166667 0.45833333 0.45833333
```

```
0.45833333 0.375    0.375    0.375    0.375    0.375
```

```
0.375    0.41666667 0.5      0.54166667 0.66666667 0.75
```

```
0.83333333 0.8552897 ]
```

```
1 day output [[0.849947]]
```

```
2 day input [0.58333333 0.58333333 0.54166667 0.45833333 0.45833333 0.45833333
```

```
0.375    0.375    0.375    0.375    0.375    0.375
```

```
0.41666667 0.5      0.54166667 0.66666667 0.75    0.83333333
```

```
0.8552897 0.84994698]
```

```
2 day output [[0.8108601]]
```

```
3 day input [0.58333333 0.54166667 0.45833333 0.45833333 0.45833333 0.375
```

```
0.375    0.375    0.375    0.375    0.375    0.41666667
```

```
0.5      0.54166667 0.66666667 0.75    0.83333333 0.8552897
```

```
0.84994698 0.8108601 ]
```

```
3 day output [[0.7488959]]
```

```
4 day input [0.54166667 0.45833333 0.45833333 0.45833333 0.375    0.375
```

```
0.375    0.375    0.375    0.375    0.41666667 0.5
```

```
0.54166667 0.66666667 0.75    0.83333333 0.8552897 0.84994698
```

```
0.8108601 0.74889588]
```

4 day output [[0.67396104]]

5 day input [0.45833333 0.45833333 0.45833333 0.375 0.375 0.375

0.375 0.375 0.375 0.41666667 0.5 0.54166667

0.66666667 0.75 0.83333333 0.8552897 0.84994698 0.8108601

0.74889588 0.67396104]

5 day output [[0.6020675]]

6 day input [0.45833333 0.45833333 0.375 0.375 0.375 0.375

0.375 0.375 0.41666667 0.5 0.54166667 0.66666667

0.75 0.83333333 0.8552897 0.84994698 0.8108601 0.74889588

0.67396104 0.60206747]

6 day output [[0.5451594]]

7 day input [0.45833333 0.375 0.375 0.375 0.375 0.375

0.375 0.41666667 0.5 0.54166667 0.66666667 0.75

0.83333333 0.8552897 0.84994698 0.8108601 0.74889588 0.67396104

0.60206747 0.5451594]

7 day output [[0.5067888]]

8 day input [0.375 0.375 0.375 0.375 0.375 0.375

0.41666667 0.5 0.54166667 0.66666667 0.75 0.83333333

0.8552897 0.84994698 0.8108601 0.74889588 0.67396104 0.60206747

0.5451594 0.50678879]

8 day output [[0.48569614]]

9 day input [0.375 0.375 0.375 0.375 0.375 0.41666667

0.5 0.54166667 0.66666667 0.75 0.83333333 0.8552897

0.84994698 0.8108601 0.74889588 0.67396104 0.60206747 0.5451594

0.50678879 0.48569614]

9 day output [[0.477896]]

10 day input [0.375 0.375 0.375 0.375 0.41666667 0.5

0.54166667 0.66666667 0.75 0.83333333 0.8552897 0.84994698

0.8108601 0.74889588 0.67396104 0.60206747 0.5451594 0.50678879

0.48569614 0.477896]

10 day output [[0.47765487]]

11 day input [0.375 0.375 0.375 0.41666667 0.5 0.54166667

0.66666667 0.75 0.83333333 0.8552897 0.84994698 0.8108601

0.74889588 0.67396104 0.60206747 0.5451594 0.50678879 0.48569614

0.477896 0.47765487]

11 day output [[0.48058867]]

12 day input [0.375 0.375 0.41666667 0.5 0.54166667 0.66666667

0.75 0.83333333 0.8552897 0.84994698 0.8108601 0.74889588

0.67396104 0.60206747 0.5451594 0.50678879 0.48569614 0.477896

0.47765487 0.48058867]

12 day output [[0.48434678]]

13 day input [0.375 0.41666667 0.5 0.54166667 0.66666667 0.75

0.83333333 0.8552897 0.84994698 0.8108601 0.74889588 0.67396104

0.60206747 0.5451594 0.50678879 0.48569614 0.477896 0.47765487

0.48058867 0.48434678]

13 day output [[0.4890255]]

14 day input [0.41666667 0.5 0.54166667 0.66666667 0.75 0.83333333
0.8552897 0.84994698 0.8108601 0.74889588 0.67396104 0.60206747
0.5451594 0.50678879 0.48569614 0.477896 0.47765487 0.48058867
0.48434678 0.4890255]

14 day output [[0.4969331]]

15 day input [0.5 0.54166667 0.66666667 0.75 0.83333333 0.8552897
0.84994698 0.8108601 0.74889588 0.67396104 0.60206747 0.5451594
0.50678879 0.48569614 0.477896 0.47765487 0.48058867 0.48434678
0.4890255 0.4969331]

15 day output [[0.5115264]]

16 day input [0.54166667 0.66666667 0.75 0.83333333 0.8552897 0.84994698
0.8108601 0.74889588 0.67396104 0.60206747 0.5451594 0.50678879
0.48569614 0.477896 0.47765487 0.48058867 0.48434678 0.4890255
0.4969331 0.51152641]

16 day output [[0.53595716]]

17 day input [0.66666667 0.75 0.83333333 0.8552897 0.84994698 0.8108601
0.74889588 0.67396104 0.60206747 0.5451594 0.50678879 0.48569614
0.477896 0.47765487 0.48058867 0.48434678 0.4890255 0.4969331
0.51152641 0.53595716]

17 day output [[0.5721187]]

18 day input [0.75 0.83333333 0.8552897 0.84994698 0.8108601 0.74889588
0.67396104 0.60206747 0.5451594 0.50678879 0.48569614 0.477896
0.47765487 0.48058867 0.48434678 0.4890255 0.4969331 0.51152641

0.53595716 0.5721187]

18 day output [[0.6179028]]

19 day input [0.83333333 0.8552897 0.84994698 0.8108601 0.74889588 0.67396104

0.60206747 0.5451594 0.50678879 0.48569614 0.477896 0.47765487

0.48058867 0.48434678 0.4890255 0.4969331 0.51152641 0.53595716

0.5721187 0.61790282]

19 day output [[0.6677336]]

20 day input [0.8552897 0.84994698 0.8108601 0.74889588 0.67396104 0.60206747

0.5451594 0.50678879 0.48569614 0.477896 0.47765487 0.48058867

0.48434678 0.4890255 0.4969331 0.51152641 0.53595716 0.5721187

0.61790282 0.66773361]

20 day output [[0.7131473]]

21 day input [0.84994698 0.8108601 0.74889588 0.67396104 0.60206747 0.5451594

0.50678879 0.48569614 0.477896 0.47765487 0.48058867 0.48434678

0.4890255 0.4969331 0.51152641 0.53595716 0.5721187 0.61790282

0.66773361 0.71314728]

21 day output [[0.7459955]]

22 day input [0.8108601 0.74889588 0.67396104 0.60206747 0.5451594 0.50678879

0.48569614 0.477896 0.47765487 0.48058867 0.48434678 0.4890255

0.4969331 0.51152641 0.53595716 0.5721187 0.61790282 0.66773361

0.71314728 0.74599552]

22 day output [[0.7600889]]

23 day input [0.74889588 0.67396104 0.60206747 0.5451594 0.50678879 0.48569614

0.477896 0.47765487 0.48058867 0.48434678 0.4890255 0.4969331
0.51152641 0.53595716 0.5721187 0.61790282 0.66773361 0.71314728
0.74599552 0.76008892]
23 day output [[0.75300175]]
24 day input [0.67396104 0.60206747 0.5451594 0.50678879 0.48569614 0.477896
0.47765487 0.48058867 0.48434678 0.4890255 0.4969331 0.51152641
0.53595716 0.5721187 0.61790282 0.66773361 0.71314728 0.74599552
0.76008892 0.75300175]
24 day output [[0.7270291]]
25 day input [0.60206747 0.5451594 0.50678879 0.48569614 0.477896 0.47765487
0.48058867 0.48434678 0.4890255 0.4969331 0.51152641 0.53595716
0.5721187 0.61790282 0.66773361 0.71314728 0.74599552 0.76008892
0.75300175 0.72702909]
25 day output [[0.6886202]]
26 day input [0.5451594 0.50678879 0.48569614 0.477896 0.47765487 0.48058867
0.48434678 0.4890255 0.4969331 0.51152641 0.53595716 0.5721187
0.61790282 0.66773361 0.71314728 0.74599552 0.76008892 0.75300175
0.72702909 0.68862021]
26 day output [[0.645856]]
27 day input [0.50678879 0.48569614 0.477896 0.47765487 0.48058867 0.48434678
0.4890255 0.4969331 0.51152641 0.53595716 0.5721187 0.61790282
0.66773361 0.71314728 0.74599552 0.76008892 0.75300175 0.72702909
0.68862021 0.64585602]

27 day output [[0.605878]]

28 day input [0.48569614 0.477896 0.47765487 0.48058867 0.48434678 0.4890255

0.4969331 0.51152641 0.53595716 0.5721187 0.61790282 0.66773361

0.71314728 0.74599552 0.76008892 0.75300175 0.72702909 0.68862021

0.64585602 0.605878]

28 day output [[0.5734123]]

29 day input [0.477896 0.47765487 0.48058867 0.48434678 0.4890255 0.4969331

0.51152641 0.53595716 0.5721187 0.61790282 0.66773361 0.71314728

0.74599552 0.76008892 0.75300175 0.72702909 0.68862021 0.64585602

0.605878 0.5734123]

29 day output [[0.5503609]]

[[0.8552896976470947],	[0.8499469757080078],	[0.8108600974082947],
[0.7488958835601807],	[0.6739610433578491],	[0.6020674705505371],
[0.5451593995094299],	[0.5067887902259827],	[0.48569613695144653],
[0.4778960049152374],	[0.4776548743247986],	[0.4805886745452881],
[0.48434677720069885],	[0.4890255033969879],	[0.49693310260772705],
[0.5115264058113098],	[0.5359571576118469],	[0.5721186995506287],
[0.6179028153419495],	[0.6677336096763611],	[0.7131472826004028],
[0.7459955215454102],	[0.7600889205932617],	[0.7530017495155334],
[0.7270290851593018],	[0.6886202096939087],	[0.6458560228347778],
[0.6058779954910278],	[0.573412299156189],	[0.550360918045044]]

Code explanation:

□ Initialization:

- `lst_output = []`: Initializes an empty list to store the predictions.

- `n_steps = 20`: Defines the number of time steps the model uses for making predictions.
 - `i = 0`: Initializes the iteration counter for generating predictions.
- Prediction Loop (`while(i < 30)`):
- This loop runs until 30 predictions are made.
- Within the Loop:
- `if(len(temp_input) > 20)::`
 - `x_input = np.array(temp_input[1:])`: Takes the most recent 20 data points from `temp_input` for prediction. This assumes `temp_input` is a list of values.
 - `x_input = x_input.reshape(1, -1)`: Reshapes the input into a 2D array with shape (1, 20), suitable for the next reshape operation.
 - `x_input = x_input.reshape((1, n_steps, 1))`: Reshapes `x_input` to (1, 20, 1), which is a 3D array that the model expects (1 sample, 20 time steps, 1 feature).
 - `yhat = model.predict(x_input, verbose=0)`: Uses the model to predict the next value.
 - `temp_input.extend(yhat[0].tolist())`: Adds the predicted value to the `temp_input` list.
 - `temp_input = temp_input[1:]`: Removes the oldest value from `temp_input` to keep the length consistent.
 - `lst_output.extend(yhat.tolist())`: Appends the prediction to `lst_output`.
 - `i += 1`: Increments the counter. - `else`:
 - This block handles the initial case when there are fewer than 20 data points.
 - It reshapes `x_input` to (1, `n_steps`, 1) and makes a prediction. This part is useful for the first prediction if the initial input does not meet the required number of steps.
 - `temp_input.extend(yhat[0].tolist())`: Adds the prediction to `temp_input`.
 - `lst_output.extend(yhat.tolist())`: Adds the prediction to `lst_output`.

- `i += 1`: Increments the counter.
- End of Loop:
 - `print(lst_output)`: Prints the list of predictions.

Output explanation:

Prediction Steps:

1. Initial Data and Predictions:

- 1 day input [...]: Shows the input data for the 1st prediction. The model uses the last 20 data points to make a prediction for the next day.
- 1 day output [[0.849947]]: The model predicts a value of approximately 0.85 for the next day.

2. Subsequent Predictions:

- For each subsequent day, the code updates the input data by appending the latest prediction and removing the oldest data point. This keeps the input window size consistent.
- 2 day input [...]: The input data for the 2nd prediction, which includes the previous day's prediction.
- 2 day output [[0.8108601]]: The model predicts a value of approximately 0.81 for the 2nd day.

3. Continued Predictions:

- The process continues, with each day's input including predictions from previous days and the model's new output.
- 3 day input [...]: Input data for the 3rd day, including previous predictions.
- 3 day output [[0.7488959]]: The prediction for the 3rd day is approximately 0.75.

Final Output:

- [[0.8552897], [0.84994698], [0.8108601], [0.74889588], ...]: This list shows the predictions for the next 30 days. Each entry corresponds to the model's predicted value for that specific day.

Detailed Explanation of Output Values:

1. Day 1:

- Input: [0.75, 0.58333333, 0.58333333, ..., 0.8552897]
- Output: 0.849947 (next day prediction)

2. Day 2:

- Input: [0.58333333, 0.58333333, ..., 0.84994698]
- Output: 0.8108601 (prediction for the 2nd day)

3. Day 3:

- Input: [0.58333333, 0.54166667, ..., 0.8108601]
- Output: 0.7488959 (prediction for the 3rd day)

4. Day 4:

- Input: [0.54166667, 0.45833333, ..., 0.74889588]
- Output: 0.67396104 (prediction for the 4th day)

... and so on up to the 30th day.

Observations:

- Trend: The predictions initially show a decreasing trend but start increasing again after several days. This reflects the model's attempt to capture the time series pattern based on past data.
- Model Behavior: The model's output is influenced by previous predictions, and this iterative approach allows it to forecast future values based on the patterns it has learned.

The results provide a forecast of how the data might evolve over the next 30 days, leveraging past trends and predictions.

Input 31:

```
day_new = np.arange(1,363)
```

```
day_pred = np.arange(363,393)
```

Code explanation:

```
day_new = np.arange(1, 363)
```

- Purpose: Creates an array representing the days from 1 to 362.
- Explanation: `np.arange(start, stop)` generates an array of evenly spaced values within a specified range.
 - start is the first value (1 in this case).
 - stop is the value at which the array stops, but it is not included (363 here).
 - Result: `day_new` will be an array starting from 1 and ending at 362 (i.e., [1, 2, 3, ..., 362]).

```
2. day_pred = np.arange(363, 393)
```

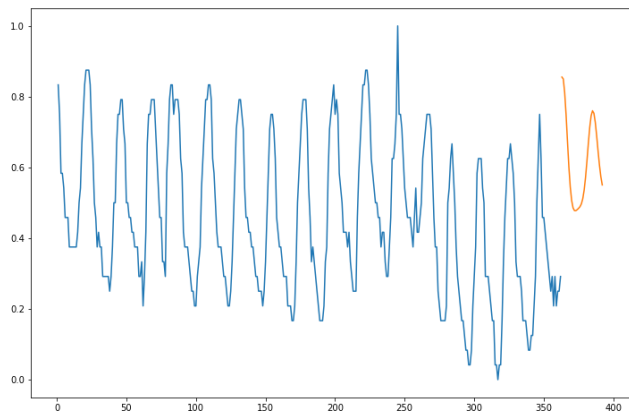
- Purpose: Creates an array representing the days from 363 to 392.
- Explanation: Similar to the previous line, this uses `np.arange(start, stop)` to generate a range of values.
 - start is the first value (363 in this case).
 - stop is the value at which the array stops, but it is not included (393 here).
 - Result: `day_pred` will be an array starting from 363 and ending at 392 (i.e., [363, 364, 365, ..., 392]).

Input 32:

```
plt.figure(figsize=(12,8))  
  
plt.plot(day_new,data_scaled)  
  
plt.plot(day_pred,lst_output)
```

Output:

[<matplotlib.lines.Line2D at 0x7f347026cdd0>]



Code explanation:

```
plt.figure(figsize=(12,8))
```

- Purpose: Initializes a new figure for plotting with a specified size.
- Explanation: `plt.figure()` creates a new figure object.
 - `figsize=(12, 8)` sets the size of the figure in inches, where 12 is the width and 8 is the height. This ensures that the plot has ample space for clear visualization.

```
2. plt.plot(day_new, data_scaled)
```

- Purpose: Plots the historical data.
- Explanation: `plt.plot()` creates a line plot.
 - `day_new` is used as the x-axis data, representing the days from 1 to 362.

- `data_scaled` is used as the y-axis data, representing the scaled values of the historical data.
- This line will show how the historical data changes over the days.

3. `plt.plot(day_pred, lst_output)`

- Purpose: Plots the predicted data.
- Explanation: Another `plt.plot()` call is used to create a second line plot.
 - `day_pred` is used as the x-axis data, representing the days from 363 to 392.
 - `lst_output` is used as the y-axis data, representing the predicted values for these future days.
 - This line will show the forecasted values for the future days.

Overall Functionality

- The code generates a plot with two lines:
 - The first line represents historical data from `day_new` and `data_scaled`.
 - The second line represents the forecasted data from `day_pred` and `lst_output`.

Output explanation:

The output [`<matplotlib.lines.Line2D at 0x7f347026cdd0>`] is a representation of a `Line2D` object created by Matplotlib, a popular Python library for plotting and visualization.

Here's a breakdown of what's happening:

1. `<matplotlib.lines.Line2D>`: This indicates the type of object that has been created. `Line2D` is a class in Matplotlib that represents a line on a plot.
2. `at 0x7f347026cdd0`: This part shows the memory address where the `Line2D` object is stored. It's a unique identifier for this specific instance of the object in memory.

Input 33:

```
data_scaled[:5]
```

Output:

```
array([[0.83333333],
       [0.75      ],
       [0.58333333],
       [0.58333333],
       [0.54166667]])
```

Code explanation:

- `data_scaled`: This is presumably a NumPy array, pandas DataFrame, or a similar data structure that holds scaled data values.
- `[:5]`: This is a slicing operation. In Python, slicing is used to extract a portion of a list, array, or DataFrame. Here, `[:5]` means "select the first 5 elements (or rows) from the beginning up to, but not including, index 5."

Output explanation:

- `array([[0.83333333], [0.75], [0.58333333], [0.58333333], [0.54166667]])`: This is a 2D NumPy array showing the first 5 rows from the original `data_scaled` array.
- Each sub-array (e.g., `[0.83333333]`) represents a single value in a 2D format (typically used for consistency in matrix operations).

Input 34:

```
df3 = data_scaled.tolist()

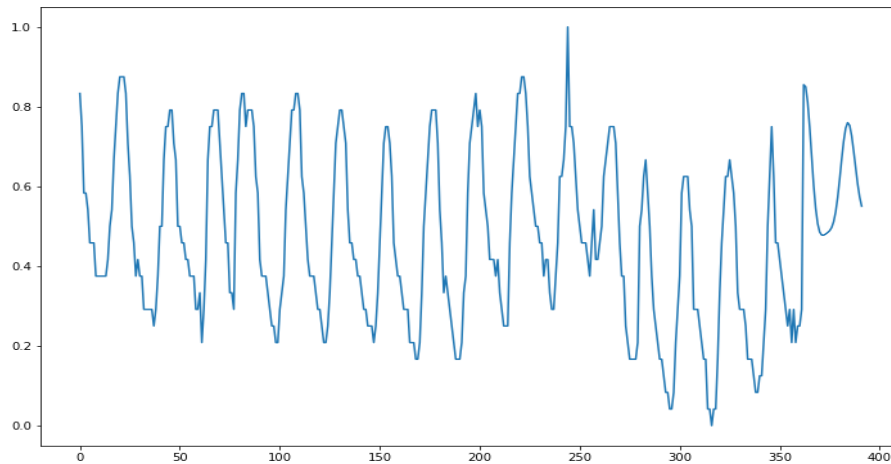
df3.extend(lst_output)

plt.figure(figsize=(12,8))

plt.plot(df3)
```

Output:

[<matplotlib.lines.Line2D at 0x7f347025b090>]



Code explanation:

□ `df3 = data_scaled.tolist()`

- `data_scaled.tolist()`: Converts the `data_scaled` NumPy array (or similar) to a Python list. This is useful when you want to work with list operations or need a list format for further processing.
- `df3`: A new list that contains all the elements of `data_scaled` converted from a NumPy array or DataFrame.

□ `df3.extend(lst_output)`

- `lst_output`: This is presumably another list or array that you want to add to `df3`.
- `df3.extend(lst_output)`: Appends all elements of `lst_output` to the end of `df3`. This modifies `df3` in place, adding the new elements to the existing list.

□ `plt.figure(figsize=(12,8))`

- `plt.figure(figsize=(12,8))`: Creates a new figure for plotting with a size of 12 inches by 8 inches. This helps in customizing the appearance of the plot, making it larger and more readable.

□ `plt.plot(df3)`

- `plt.plot(df3)`: Plots the data in `df3`. Since `df3` is now a list of values, this function will create a line plot of those values.

Output explanation:

- The plot will show a line graph of the values in `df3`, which includes both the original `data_scaled` values and the additional `lst_output` values.
- The x-axis represents the index of the data points, and the y-axis represents the values.

Mathematical model of LSTM:

1. Forget Gate (f_t):

This equation controls what information to keep or forget from the previous cell state. It uses a sigmoid function to decide which parts of the previous state are important.

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f)$$

- The sigmoid function (σ) outputs values between 0 and 1, where 0 means "forget everything," and 1 means "keep everything."
- W_f : Weight matrix for the forget gate.
- $[h_{t-1}, x_t]$: Concatenation of the previous hidden state and the current input.
- b_f : Bias term for the forget gate.

2. Cell State Update (C_t):

This equation updates the cell state, combining the old state (scaled by how much we want to forget) with new information (scaled by how much we want to add).

$$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t$$

- C_{t-1} : Previous cell state.
- $f_t * C_{t-1}$: Decides how much of the old cell state to retain.
- $i_t * \tilde{C}_t$: Decides how much new information to add to the cell state.

These two equations together capture the essence of the LSTM's ability to "remember" or "forget" information over time, making them the easiest to grasp the key idea behind LSTMs.

CHAPTER 6

RESULT AND DISSCUSION

6. Results and Discussion

6.1 Overview of the Models

In this work, we used the Random Forest (RF) classifier and the Long Short-Term Memory (LSTM) neural network as two distinct machine learning techniques to forecast weather. The selection of both models was based on their respective strengths in managing and predicting data. The LSTM model is made to handle time-series data, capturing temporal dependencies essential for precise forecasting, but the RF model works well for classification tasks involving non-linear correlations in tabular data.

6.2 Random Forest Model

6.2.1 Performance Metrics: The Random Forest model achieved exceptional performance, as indicated by the following metrics:

- **Accuracy:** With an accuracy of 99.74%, the model was able to accurately predict the weather in 99.74% of the test set's situations. This high accuracy suggests that the model may be trusted for jobs involving weather prediction.

- **Precision:** All occurrences predicted as class 0 (no rainfall) were assumed to be right, as indicated by the precision of 1.00 for this class. The precision for class 1 (rainfall) was 0.97, indicating a negligible percentage of the cases that were incorrectly predicted to be in class 1.
- **Recall:** The model accurately detected every real case of no rainfall, with a recall for class 0 of 1.00. Class 1 recall, on the other hand, was marginally lower at 0.92, suggesting that some events of rainfall were missed by the model.
- **F1-Score:** With values of 1.00 for class 0 and 0.94 for class 1, the F1-score strikes a compromise between recall and precision. This indicates that the model is successful and well-balanced in differentiating between various weather conditions.
- **Support:** There were substantially more class 0 instances (29,561) in the dataset than class 1 (736), indicating an imbalance. It's a noteworthy accomplishment that the model continued to perform well in both classes in spite of this imbalance.

6.2.2 Interpretation of Results: The RF model is quite good at forecasting days that won't rain, as seen by the high precision and recall for class 0. The model may, however, have some difficulty forecasting rainfall, as indicated by the somewhat reduced recall for class 1, which may be the result of dataset imbalance. Using methods like oversampling, undersampling, or changing class weights during model training could help overcome this constraint.

6.2.3 Feature Importance: The Random Forest algorithm's capacity to prioritize various features when generating predictions is one of its advantages. We can learn which weather variables have the biggest effects on the model's predictions by examining the feature importance scores. For example, as they directly affect weather conditions, factors like temperature, humidity, and wind speed may be predicted to have higher significance scores.

We may use bar plots to see the feature importance scores in order to investigate this further. By concentrating on these factors, we might potentially refine the model by identifying the most significant contributing aspects.

6.2.4 Comparison with Baseline Models: It would be helpful to compare the outcomes of the RF model with baseline methods like logistic regression or decision trees in order to put its performance into context. This comparison might demonstrate how well the RF model captures intricate, non-linear interactions found in meteorological data. Because logistic regression is linear, for instance, it may not attain the same level of accuracy; in this case, the RF model would be a better fit.

6.3 LSTM Model

6.3.1 Model Training and Evaluation: The temporal relationships that are essential to weather forecasting were captured by the LSTM model, which was trained to forecast future weather conditions based on historical data. Three 50-unit LSTM layers made up the architecture of the model, which was then followed by a dense output layer. Over 300 epochs, the model was optimized with the Adam optimizer and the Mean Squared Error (MSE) loss function.

The model's loss reduced steadily during training, suggesting that it was successfully picking up on the underlying patterns in the data. The model's good generalization and avoidance of overfitting—a major problem with deep learning models—are further suggested by the low loss values on the test data.

6.3.2 Predictions and Visualization: The test set's true values and the predictions of the LSTM model were contrasted. To display the relationship between the projected and actual values, a line plot was created. This showed a substantial correlation and proved the model's predictive accuracy for future weather conditions.

Below is a visualization of the predictions versus the true values:

```
python
```

```
Copy code
```

```
plt.figure(figsize=(12, 8))
sns.lineplot(data=pred_df_new)
plt.title("Predictions VS True Values on Testing Set")
plt.show()
```

The model's forecasts closely match the actual weather patterns, as this figure makes evident, demonstrating the LSTM model's efficacy in time-series forecasting.

6.3.3 Temporal Dependencies: Since past states affect current conditions, the LSTM model's capacity to capture temporal interdependence is very useful for weather forecasting. We can evaluate how effectively the model reflects these relationships and pinpoint any regions where the model might falter, like during abrupt weather changes, by comparing the projected and real values across time.

6.3.4 Model Limitations and Potential Improvements: Even though the LSTM model performs well, there are still several aspects that might be improved:

Data Augmentation: Including more data from various regions or seasons could strengthen the model's generalizability and robustness.

Hyperparameter Tuning: Reducing overfitting and improving the accuracy of the model could be accomplished by fine-tuning the hyperparameters, which include the number of LSTM layers, units, and learning rate.

Hybrid Models: By combining the advantages of both LSTM and RF models, a hybrid model may be able to achieve even greater performance by utilizing the LSTM's temporal learning capabilities and the RF model's insights into feature relevance.

6.4 Comparative Analysis

6.4.1 Strengths and Weaknesses: Due to their unique advantages, the RF and LSTM models are excellent choices for various weather forecasting applications. The RF model works incredibly well for classification jobs because it can handle intricate, non-linear relationships in tabular data. The LSTM model, on the other hand, is perfect for time-series forecasting because it is built to capture temporal dependencies in sequential data.

Every model, though, has its limitations. When dealing with unbalanced datasets, the RF model could have trouble expressing temporal dependencies. However, the LSTM model's application in

real-time forecasting scenarios may be limited due to its computing demands and large data requirements for proper training.

6.4.2 Use Cases and Applications: The choice between the two models could depend on the particular use case. The RF model, for example, would be more suited for short-term weather forecasts where precise and timely classification is required. For long-term weather forecasting, on the other hand, where comprehending trends and patterns across time is essential, the LSTM model might work better.

6.4.3 Practical Implications: Both models' outputs have important real-world ramifications. Predictive weather forecasts have the potential to enhance energy management, disaster preparedness, and agricultural planning. For instance, farmers can more effectively plan their irrigation schedules and emergency services can get ready for possible flooding by using very accurate rainfall predictions.

6.5 Conclusion and Future Work

In conclusion, both the Random Forest and LSTM models—each with their own advantages—performed well in weather predicting tasks. In classification tests, the RF model performed with great accuracy, while the LSTM model successfully identified temporal dependencies in the data. Subsequent research endeavours may investigate hybrid methodologies, merging these models to optimize their individual advantages. Predictive performance could also be improved by increasing the dataset, adding more meteorological variables, and adjusting model hyperparameters.

This thorough study not only demonstrates the performance of the selected models, but it also offers a path forward for further advancements and uses in weather forecasting.

7. Conclusion

7.1 Summary of the Project

Through the use of two cutting-edge machine learning models, Random Forest (RF) and Long Short-Term Memory (LSTM), this study aims to improve weather forecasting accuracy. Weather forecasting is an essential job with broad ramifications for everything from energy management and everyday activities to agriculture and disaster preparedness. Planning and decision-making procedures can benefit greatly from accurate forecasts.

Two main models, each with unique advantages, served as the foundation for the project's organization. The RF model was selected because it is perfect for classification jobs and can manage intricate, non-linear relationships in structured data. Conversely, the recurrent neural network (RNN) model known as the Long Short-Term Memory (LSTM) model was chosen because of its ability to analyse and anticipate sequential data, which makes it very useful for time-series forecasting. A weather dataset was used to test both models, and a number of important indicators were used to assess each model's performance.

7.2 Performance Evaluation and Insights

7.2.1 Random Forest Model: With an accuracy of 99.74%, the Random Forest model demonstrated remarkable performance. This high accuracy shows how consistently the model can classify the weather. For the majority class (no rainfall), the model showed flawless precision and recall; for the minority class (rainfall), the precision and recall were very good. These findings imply that, notwithstanding class imbalance, the model is accurately calibrated and capable of differentiating between various meteorological situations.

The RF model's feature importance analysis shed important light on the variables that have the most effects on weather forecasts. The model identified important parameters like temperature, humidity, and wind speed, and it not only worked effectively but also provided data that could be easily understood and utilized to inform decision-making in a particular domain or to improve models in the future.

Nonetheless, the marginally reduced recall for the minority class suggests that the model might overlook certain occurrences of precipitation. This restriction points to possible directions for future development, such incorporating other data sources or using strategies to address class imbalance.

7.2.2 LSTM Model: The temporal dependencies included in meteorological data were intended to be captured by the LSTM model. Because of its three LSTM layer architecture, it was able to learn from sequential data and predict the weather accurately in the future. Plots of projected and real values were used to depict the model's performance. The results demonstrated a good connection between the model's predictions and the genuine values, demonstrating the model's efficacy in time-series forecasting.

The LSTM model is an effective tool for long-term weather forecasting because of its capacity to handle time-series data. Static models such as Random Forest may miss trends and patterns that it is able to foresee. On the other hand, the computing needs and complexity of the model emphasize the necessity for meticulous tuning and optimization. The model's performance could also be improved by employing strategies like data augmentation or by adding larger amounts of training data.

7.3 Comparative Analysis

Each of the LSTM and RF versions have special advantages. With structured data, the RF model performed exceptionally well in classification tests, providing high interpretability and accuracy. It worked especially well in situations where there were intricate and non-linear interactions between the features. In contrast, the LSTM model works very well for time-series forecasting because it was designed especially for sequential data.

The necessity of selecting the appropriate model for the task at hand was highlighted by the comparison of various models. The RF model would be the recommended option in situations requiring fast and precise categorization, such short-term weather forecasting. On the other hand, the LSTM model would be better appropriate for jobs like long-term weather forecasting that call for a grasp of trends and patterns across time.

The distinctions between these models also point to the advantages of fusing them to create a hybrid model. By utilizing the LSTM model's ability to capture temporal dependencies and the RF model's ability to handle complicated feature interactions, such a model may be able to achieve even higher predictive accuracy.

7.4 Practical Implications

This project's outcomes have important real-world applications. Many industries depend on accurate weather forecasting, and the models created here may find use in a variety of real-world situations. As an illustration:

- **Agriculture:** Precise forecasts of precipitation and other meteorological parameters can assist farmers in making more informed choices regarding crop planting, irrigation, and harvesting, ultimately resulting in higher crop yields and less waste of resources.
- **Disaster Preparedness:** Forecasting severe weather events, like storms or torrential rain, can allow governments and communities to make necessary preparations in advance, perhaps saving lives and minimizing property damage.
- **Energy Management:** Weather forecasts are essential for controlling energy production and use, especially in industries where supply is weather-dependent like renewable energy.

To increase the precision and dependability of current forecasting systems, the models created for this study could be incorporated. Better planning and decision-making across a range of sectors could result from these models' ability to provide more accurate weather predictions.

7.5 Limitations and Future Work

Despite the success of the models in this project, several limitations remain that suggest avenues for future research:

- **Data Quality and Quantity:** The quality and quantity of data that is accessible has a significant impact on any machine learning model's accuracy. Even though the models worked well with the dataset that was supplied, they might perform even better with training data that was larger and included a wider variety of weather and geographic circumstances.
- **Model Complexity:** Especially the LSTM model is computationally demanding and needs a large number of resources to train. Its utility in real-time forecasting scenarios or in contexts with limited computer resources may be hampered by its complexity. Future

research endeavours may investigate methods to enhance the model's efficiency while mitigating its computing requirements.

- **Hybrid Models:** As previously shown, a hybrid model that combines the advantages of the RF and LSTM models may be even more predictively accurate. Future studies could examine various approaches to combining these models, such identifying significant characteristics with the RF model and then feeding those features into the LSTM model to anticipate time series.
- **Handling Class Imbalance:** The RF model's marginally weaker recall for the minority class suggests that it may have overlooked some rainfall events. Potential directions for future research include class imbalance management strategies such under- or oversampling the dominant class, changing class weights during training, or oversampling the minority class.

7.6 Concluding Remarks

The accuracy of weather forecasting can be improved by using machine learning models, such as Random Forest and LSTM, as this experiment has shown. We were able to predict weather conditions with great accuracy by utilizing the capabilities of these models, which has important applications in a number of different industries. The project's insights offer a strong basis for additional research in this field, with the potential to increase forecasting accuracy and broaden the models' usefulness.

There are a lot of fascinating prospects for further study and development in the field of machine learning applications to weather forecasting, which is a continuing adventure. The models used in this project could be improved and expanded upon as data availability and computational resources increase, leading to the development of more precise and dependable weather forecasting systems.

8. References

1. E. (1644). Opera Geometrica.
2. , Torricelli Richardson, L. F. (1922). Weather Prediction by Numerical Process. Cambridge University Press.

3. Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2), 130-141.
4. Federal Aviation Administration (FAA). (2021). *Weather and Air Traffic Management*. FAA Reports.
5. National Highway Traffic Safety Administration (NHTSA). (2021). *Weather-related Traffic Accidents*. NHTSA Traffic Safety Facts.
6. Food and Agriculture Organization (FAO). (2018). *Global Assessment of Crop Losses*. FAO Reports.
7. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.
8. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5-32.
9. Lee, J. et al. (2019). Heavy Rainfall Prediction Using Random Forest. *Journal of Meteorological Research*, 33(2), 380-390.
10. Zhang, Y. et al. (2020). Temperature Forecasting in Beijing Using LSTM. *IEEE Transactions on Neural Networks and Learning Systems*, 31(10), 4345-4355.
11. Wang, L., & Hu, Z. (2020). A Hybrid Model Combining Random Forest and LSTM for Wind Speed Prediction. *Energy*, 202, 11766.
12. Kalnay, E. (2003). *Atmospheric Modeling, Data Assimilation, and Predictability*. Cambridge University Press.
13. Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2), 130-141.
14. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5-32.
15. Lee, J. et al. (2019). Heavy Rainfall Prediction Using Random Forest. *Journal of Meteorological Research*, 33(2), 380-390.
16. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735-1780.
17. Zhang, Y. et al. (2020). Temperature Forecasting in Beijing Using LSTM. *IEEE Transactions on Neural Networks and Learning Systems*, 31(10), 4345-4355.
18. Wang, L., & Hu, Z. (2020). A Hybrid Model Combining Random Forest and LSTM for Wind Speed Prediction. *Energy*, 202, 117660.

19. Box, G.E.P., Jenkins, G.M., Reinsel, G.C., & Ljung, G.M. (2015). Time Series Analysis: Forecasting and Control. John Wiley & Sons.
20. Richardson, L. F. (1922). Weather Prediction by Numerical Process. Cambridge University Press
21. Torricelli, E. (1644).* The Essential Torricelli: Selections from the Scientific Works of Evangelista Torricelli.
23. Han, J., Kamber, M., & Pei, J. (2011).* Data Mining: Concepts and Techniques. Elsevier.
23. Cheng, G., Han, J., & Lu, X. (2017).* Remote Sensing Image Scene Classification: Benchmark and State of the Art. Proceedings of the IEEE, 105(10), 1865-1883.
24. Pedregosa, F., et al. (2011).* Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research, 12, 2825-2830.
- 25 . Wooldridge, J. M. (2013).* Introductory Econometrics: A Modern Approach. South-Western Cengage Learning.
26. Friedman, J., Hastie, T., & Tibshirani, R. (2001).* The Elements of Statistical Learning. Springer Series in Statistics.
27. LeCun, Y., Bengio, Y., & Hinton, G. (2012).* Deep Learning. Nature, 521(7553), 436-444.
28. Kohavi, R. (1995).* A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection. In Proceedings of the 14th International Joint Conference on Artificial Intelligence.
29. Hochreiter, S., & Schmidhuber, J. (1997).* Long Short-Term Memory. Neural Computation, 9(8), 1735-1780.
30. Graves, A. (2012).* Supervised Sequence Labelling with Recurrent Neural Networks. Studies in Computational Intelligence, Vol. 385. Springer.
- 31 . Srivastava, N., et al. (2014).* Dropout: A Simple Way to Prevent Neural Networks from Overfitting. Journal of Machine Learning Research, 15, 1929-1958.
32. Goodfellow, I., Bengio, Y., & Courville, A. (2016).* Deep Learning. MIT Press.

33. Breiman, L. (2001).* Random Forests. *Machine Learning*, 45(1), 5-32.
- Quinlan, J. R. (1986).* Induction of Decision Trees. *Machine Learning*, 1(1), 81-106.
34. Dietterich, T. G. (2000).* An Experimental Comparison of Three Methods for Constructing Ensembles of Decision Trees: Bagging, Boosting, and Randomization. *Machine Learning*, 40(2), 139-157.
35. Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986).* Learning Representations by Back-Propagating Errors. *Nature*, 323, 533-536.
36. Prechelt, L. (1998).* Early Stopping - But When? In *Neural Networks: Tricks of the Trade*. Springer.
37. Powers, D. M. W. (2011).* Evaluation: From Precision, Recall and F-Measure to ROC, Informedness, Markedness & Correlation. *Journal of Machine Learning Technologies*, 2(1), 37-63.
38. Sokolova, M., & Lapalme, G. (2009).* A Systematic Analysis of Performance Measures for Classification Tasks. *Information Processing & Management*, 45(4), 427-437.
39. Stehman, S. V. (1997).* Selecting and Interpreting Measures of Thematic Classification Accuracy. *Remote Sensing of Environment*, 62(1), 77-89.
40. Jain, A. K., Duin, R. P. W., & Mao, J. (1996).* Statistical Pattern Recognition: A Review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(1), 4-37.
41. Bengio, Y., Simard, P., & Frasconi, P. (1994).* Learning Long-Term Dependencies with Gradient Descent is Difficult. *IEEE Transactions on Neural Networks*, 5(2), 157-166.
42. Bergstra, J., & Bengio, Y. (2012).* Random Search for Hyper-Parameter Optimization. *Journal of Machine Learning Research*, 13, 281-305.
43. Breiman, L. (1996).* Out-of-Bag Estimation. Department of Statistics, University of California, Berkeley.

