

Implementing Weather Prediction Using Physics Informed Neural Networks (PINNS)

¹P. BINDHU PRIYA, ²MANDAVAKURITI HARIKA

¹Assistant Professor, ²MCA Final Semester,

Master of Computer Applications,

Sanketika Vidya Parishad Engineering College, Vishakhapatnam, Andhra Pradesh, India

ABSTRACT

Physics-based neural networks (pinn's) have proven to be a promising methodology for combining domain knowledge and data control learning, particularly in modeling complex dynamic systems. This study presents a hybrid deep learning framework that integrates physics-based limitations for predicting climate variables and Bidirectional Long Short-Term Memory (BiLSTM). The aim is to predict atmospheric conditions near the creation, particularly temperature and geopotential heights, using continuous observations from the created data. The BiLSTM model is trained to simultaneously record the underlying time patterns of data and is in compliance with the physics of atmospheric processes. The concept of physical loss is introduced. This is derived from the simplified thermal diffusion equation to punish violations of the basic energy diffusion properties. This loss of physics, combined with the standard loss of standard square error (MSE), ensures that the model's predictions are not only accurate but physically consistent. Comparative reviews show that physical models improve prediction stability and achieve greater compliance with physical principles compared to purely data-driven baselines. Furthermore, the addition of physics-based regularization improves generalization through invisible samples, which helps reduce overadaptation, especially in border regions where traditional models often fail. By embedding physical knowledge directly into the training process, this model provides a way to reliable and interpretive weather and climate prediction systems that have a more comprehensive effect on promoting scientific machine learning in modeling the Earth system.

IndexTerms: Physics-Informed Neural Networks (PINNs), Bidirectional Long Short-Term Memory (BiLSTM), Climate Forecasting, Geopotential Height, Temperature Prediction, Diffusion Equation, Physics-Based Loss, Time Series Forecasting.

1.INTRODUCTION

Physics-Informed Neural Networks (PINNs) integrate physical laws into the training process of deep neural networks, enabling the models to make predictions that conform to known governing equations.[5] This makes them especially useful in scientific fields such as climate modeling, where it is crucial to maintain physical consistency for accurate forecasts. This project employs a hybrid model that combines Bidirectional Long Short-Term Memory (BiLSTM) networks with physics-informed loss functions to enhance the prediction of atmospheric variables like temperature and geopotential height. The climate forecasting system is built upon time-series data derived from reanalysis datasets, focusing on temperature fields and geopotential height distributions.[9] Each location within the climate grid is analyzed individually, using historical time sequences as input to the deep learning model. The BiLSTM architecture is capable of capturing sequential dependencies in both forward and backward directions, offering a more complete view of atmospheric dynamics over time. To improve the realism of the simulation, a physics-based loss function, grounded in the heat diffusion equation, is added. This ensures that predicted temperature fields conform to basic physical laws. By incorporating these physical principles, the model not only increases learning efficiency and reduces

overfitting but also improves generalization in predicting future climate states. The resulting PINN–BiLSTM model provides climate scientists and meteorologists with a robust, interpretable, and physically consistent weather prediction tool.[13]

1.1 Existing System

The existing systems for climate forecasting primarily rely on traditional numerical weather prediction (NWP) models and pure data-driven deep learning models.[16] While NWP models are physically accurate, they are computationally expensive, require massive resources, and often struggle with high-resolution global simulations. On the other hand, pure deep learning models such as standard LSTM or CNNs are efficient but lack physical interpretability, often generating outputs that violate known scientific laws. These models treat data as independent sequences and ignore the underlying physics that govern atmospheric dynamics. As a result, they may perform poorly when generalizing to new or edge-case conditions.[20] Additionally, they require large labeled datasets, are prone to overfitting, and may fail to produce reliable predictions in data-scarce environments. Most models also do not incorporate conservation laws, such as those governing heat diffusion, leading to inaccurate or unrealistic forecasts. Current approaches also lack robust mechanisms to ensure consistency across spatial and temporal dimensions. Moreover, the absence of physically-informed regularization reduces model stability. Consequently, they cannot fully address real-world challenges like long-term forecasting, uncertainty quantification, and regional variability. These limitations highlight the need for models that balance data-driven learning with physical constraints to enhance reliability. The existing system, though advanced, still struggles to deliver physically consistent, efficient, and interpretable weather prediction at scale.[14]

1.1.1 Challenges:

- **Lack of Physical Consistency**

Purely data-driven models often ignore physical laws (like heat diffusion), leading to physically unrealistic predictions.[8]

- **High Computational Cost in Traditional Models**

Numerical weather prediction (NWP) models are accurate but require heavy computational resources, making them impractical for real-time or large-scale deployment.

- **Poor Generalization**

Standard deep learning models often overfit to training data and struggle to generalize to unseen spatial-temporal scenarios.

- **Handling Noisy or Incomplete Data**

These models perform poorly when working with sparse, incomplete, or noisy datasets, which are common in real-world atmospheric data.

- **Sequential Limitations**

Unidirectional models fail to capture backward temporal dependencies, leading to incomplete understanding of atmospheric time-series.

- **Limited Scalability Across Regions**

Existing methods don't generalize well across different geographic locations due to lack of integrated physical reasoning.[2]

1.2 Proposed system:

The proposed system introduces a hybrid deep learning model that integrates Physics-Informed Neural Networks (PINNs) with Bidirectional Long Short-Term Memory (BiLSTM) networks to enhance climate forecasting accuracy and reliability.[10] Unlike traditional models, this system embeds physical laws—particularly the heat diffusion equation—directly into the training process through a physics-based loss function. It utilizes time-series data from reanalysis datasets, focusing on temperature and geopotential height

across a spatial climate grid.[19] The BiLSTM architecture captures both past and future dependencies in atmospheric patterns, ensuring a more complete temporal understanding. The PINN component enforces physical consistency, helping the model generate scientifically valid outputs even when data is sparse or noisy. The use of diffusion-based regularization improves stability and generalization, especially in edge regions where traditional models often fail. The model is implemented in PyTorch, and training involves minimizing a combined loss: mean squared error and physics-informed loss. A neighbor table is constructed to compute the Laplacian efficiently, ensuring the model adheres to physical diffusion behavior. The system demonstrates high accuracy, low error rates, and robust generalization to unseen data. Visualizations such as scatter plots and heatmaps confirm its predictive precision. By blending data-driven learning with physical principles, the proposed model outperforms purely statistical or computational models. It is scalable, interpretable, and suitable for real-world deployment in weather prediction tasks. Overall, the proposed system provides a scientifically grounded, efficient, and accurate framework for climate forecasting.[7]

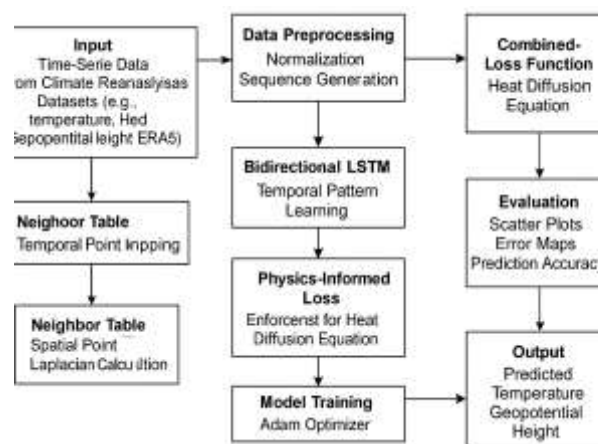


Fig: 1 Proposed Diagram

1.2.1 Advantages:

- **Physically Consistent Predictions**

Integrating the heat diffusion equation ensures that predictions align with fundamental physical laws, enhancing realism.[11]

- **Improved Accuracy**

The combination of PINNs with BiLSTM improves the model's ability to accurately forecast temperature and geopotential height.

- **Captures Temporal Dependencies**

BiLSTM architecture learns both past and future time dependencies, improving the understanding of sequential climate patterns.

- **Better Generalization**

Physics-based regularization reduces overfitting, allowing the model to perform well even on unseen or noisy data.[17]

2.1 Architecture:

1.Input Layer

The input to the system comes from reanalysis climate datasets like ERA5. These datasets include important atmospheric variables such as temperature and geopotential height collected over time. Each data point represents climate information at specific grid locations across various time steps. This time-series data serves as the foundation for training the prediction model.[4]

2. Data Preprocessing

Before feeding the data into the model, it goes through a preprocessing stage. In this step, the temperature and geopotential height values are normalized to ensure consistency during training. The data is then converted

into sequences, where each sequence consists of 20-time steps, and the model learns to predict the next step. This prepares the input in a format suitable for time-series analysis.

3. Neighbours Table Construction

For each grid point (location) in the dataset, the system calculates its neighboring points — left, right, top, and bottom. This is important for applying the Laplacian operator, which is needed to calculate the physics-based loss. The neighbor table helps identify how data at each point is related to its surrounding area, which supports physical consistency in predictions.

4. Bidirectional LSTM (BiLSTM)

The model uses a Bidirectional Long Short-Term Memory (BiLSTM) network to learn from the time-series data. This deep learning structure captures patterns in both forward and backward directions of time. That means the model not only considers past data but also future trends within the sequence, which improves prediction accuracy over time.[12]

5. Physics-Informed Loss Function

To make the model's predictions scientifically accurate, a physics-based loss function is included. This loss is based on the heat diffusion equation, which governs how temperature spreads over time. If the model makes a prediction that doesn't follow this physical law, it is penalized, which helps guide the model to learn patterns that obey real-world physics.[6]

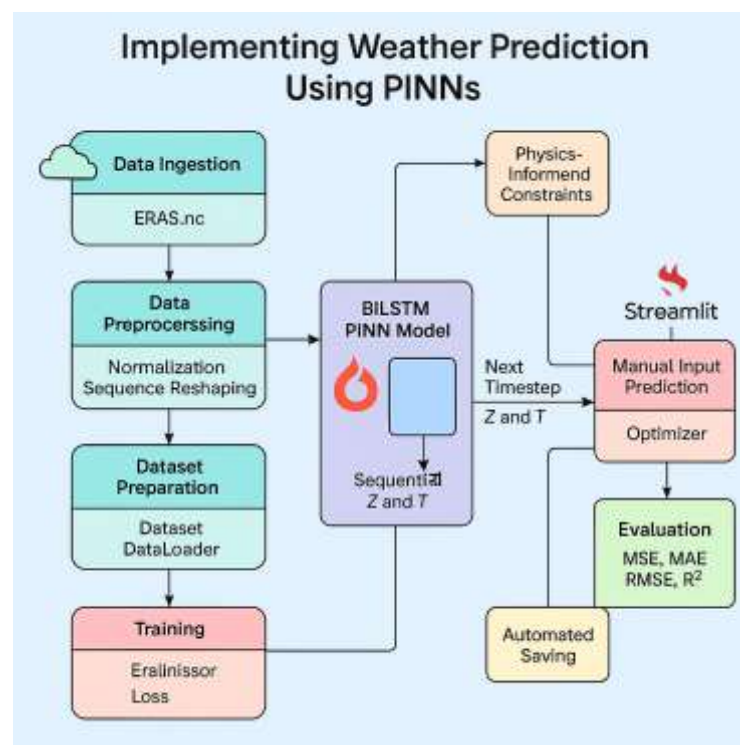


Fig:2 Architecture

UML DIAGRAMS




Fig:3 use case diagram

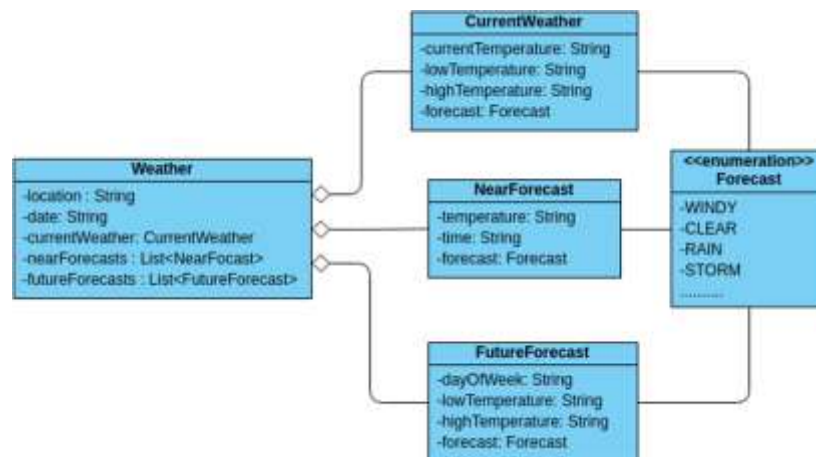


Fig:4 class diagram

2.2 Algorithm:

1. Bidirectional Long Short-Term Memory (BiLSTM)

This is a type of Recurrent Neural Network (RNN) that processes time-series data in both forward and backward directions. In this project, BiLSTM captures temporal dependencies in atmospheric data, such as temperature and geopotential height, enabling better sequence modeling for weather prediction.[18]

2. Physics-Informed Neural Networks (PINNs)

PINNs incorporate physical laws (like partial differential equations) into the learning process. In this project, a diffusion-based physical loss derived from the heat diffusion equation is introduced. It penalizes physically inconsistent predictions, ensuring the model remains scientifically accurate, not just statistically fit.

3. Mean Squared Error (MSE) Loss

A standard loss function used in regression problems. It computes the average squared difference between predicted and actual values. It forms one part of the total loss function used during model training [15].

4. Physics-Based Loss Function

This custom loss is calculated from the discrepancy between predicted and physically expected behavior, based on the Laplacian (second spatial derivative) of the temperature field. It complements MSE loss to enforce physical consistency.[3]

2.3 Techniques:

In this project, several important techniques were used to build a reliable and scientifically accurate weather prediction model. The Bidirectional LSTM (BiLSTM) helped the model learn from both past and future weather data, improving the accuracy of time-series predictions. To ensure the predictions followed real-world physical laws, the project used Physics-Informed Neural Networks (PINNs), which integrated knowledge from the heat diffusion equation into the learning process. A physics-based loss function was added alongside the traditional Mean Squared Error (MSE) loss to penalize physically incorrect predictions, helping the model stay consistent with natural behavior. The Adam optimizer was used to train the model efficiently by adjusting the learning rate automatically. To prepare the data, normalization was applied so all features were on a similar scale, and a sliding window technique was used to create smaller time-sequences (20-time steps) for easier model training. The project also built a neighbor table to identify surrounding grid points, which was important for calculating the Laplacian—a mathematical method that estimated how temperature changes across space. This supported the physics-based loss function. Lastly, various visualization tools like heatmaps, scatter plots, and loss curves were used to monitor model performance and understand prediction behavior. Together, these techniques allowed the model to produce accurate, stable, and physically meaningful weather forecasts.[1]

2.4 Tools:

- **Python 3.8:**

Python was the main programming language used to write the code for the entire project. It provided flexibility and support for scientific computing, machine learning, and data handling.

- **PyTorch:**

PyTorch was used to build and train the BiLSTM neural network model. It allowed easy integration of custom loss functions, such as the physics-informed loss used in this project.

- **NumPy:**

NumPy handled numerical operations like array transformations and statistical calculations. It helped in data preprocessing before feeding it into the model.

- **Xarray:**

Xarray was used to read and manage the multidimensional ERA5 climate data stored in NetCDF files. It simplified handling data with multiple time and spatial dimensions.

- **Torch.utils.data:**

This PyTorch utility was used to create custom datasets and data loaders. It helped in efficiently managing and feeding data batches into the model during training.

- **Scikit-learn:**

Scikit-learn was used to split the dataset and calculate evaluation metrics like RMSE, MAE, R^2 score, and ROC-AUC. It provided tools to measure how well the model performed.

- **Matplotlib:**

Matplotlib was used to plot graphs and visualizations, such as scatter plots, loss curves, and heatmaps. These visual tools helped evaluate and understand the model's behavior.

- **Jupyter Notebook / VS Code:**

Jupyter Notebook and Visual Studio Code were the coding environments used. They allowed interactive development, testing, and visualization of model outputs.

2.5 Methods:

1. Data Collection and Preprocessing

The project used atmospheric data (temperature and geopotential height) from the ERA5 reanalysis dataset. The data was normalized (mean subtraction and standard deviation scaling) to ensure stability during training.

2. Sliding Window Sequence Creation

Time-series data was split into sequences of 20-time steps, with the next time step used as the target. This method allowed the model to learn how weather variables change over time.

3. Neighbor Table Construction

For each point on the spatial grid (latitude and longitude), neighboring points (top, bottom, left, right) were identified. These were used to compute the Laplacian during physics loss calculation.

4. Model Building using BiLSTM

A Bidirectional Long Short-Term Memory (BiLSTM) model was constructed with two layers. It captured both forward and backward temporal patterns for better climate forecasting.

5. Physics-Informed Loss Integration

A custom physics-based loss was created using the heat diffusion equation. This was combined with the standard Mean Squared Error (MSE) loss to ensure predictions followed real-world physics.

3. METHODOLOGY

3.1 Input:

The input for this project is climate data taken from the ERA5 reanalysis dataset, specifically focusing on two key atmospheric variables: temperature (T) and geopotential height (Z). This data is structured as a time-series over a global grid of latitude and longitude points. Each data point represents how these variables change over time at specific locations. Before feeding into the model, the data is normalized (by subtracting the mean and dividing by the standard deviation) to ensure consistency and improve training performance. The input is then divided into sequences of 20-time steps, where each sequence helps the model learn patterns and predict the next time step values.

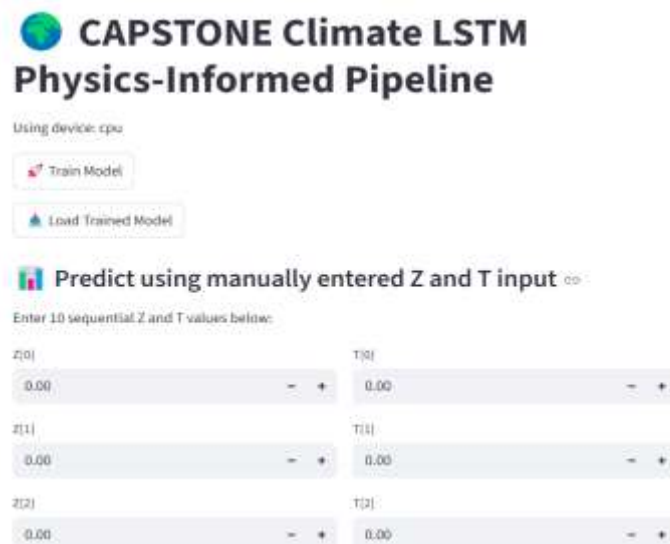


Fig:input values

3.2 Method of Process:

The method of process in this project follows a step-by-step approach to build a weather prediction model using Physics-Informed Neural Networks (PINNs) and BiLSTM. First, climate data (temperature and geopotential height) is collected from the ERA5 reanalysis dataset. This data is then preprocessed by normalizing the values and dividing it into time-series sequences using a sliding window method (20-time steps each). Next, a neighbor table is constructed for each spatial point to identify adjacent grid points, which are used later for calculating the Laplacian in the physics-based loss. The BiLSTM model is then built to learn both forward and backward time dependencies in the data. During training, a combined loss function is used—one part is Mean Squared Error (MSE) to ensure accuracy, and the other is a physics-informed loss based on the heat diffusion equation to ensure physical correctness. The model is trained using the Adam optimizer for 30 epochs. After training, the model is evaluated using performance metrics (like RMSE, MAE, R^2) and visual tools (like scatter plots, heatmaps, and loss curves) to verify accuracy and consistency.

3.3 Output:

The output of this project is the predicted values of atmospheric variables—specifically temperature and geopotential height—for future time steps based on past climate data. The model produces these predictions at specific spatial grid points and time intervals. The output is evaluated using statistical metrics such as Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and R^2 score, which indicate how closely the predictions match actual values. Additionally, visual outputs such as heatmaps, scatter plots, and loss curves help demonstrate how accurately the model has learned and how well it generalizes to unseen data. The final output confirms that the model not only predicts accurately but also maintains physical consistency by following the heat diffusion law.

Predict using manually entered Z and T input

Enter 10 sequential Z and T values below:

Z[0]	12.00	T[0]	0.00
Z[1]	5.00	T[1]	13.00
Z[2]	7.00	T[2]	1.0
Z[3]	0.00	T[3]	0.00
Z[4]	0.00	T[4]	0.00
Z[5]	0.00	T[5]	0.00
Z[6]		T[6]	
Z[7]		T[7]	
Z[8]		T[8]	
Z[9]		T[9]	

Fig:prediction values

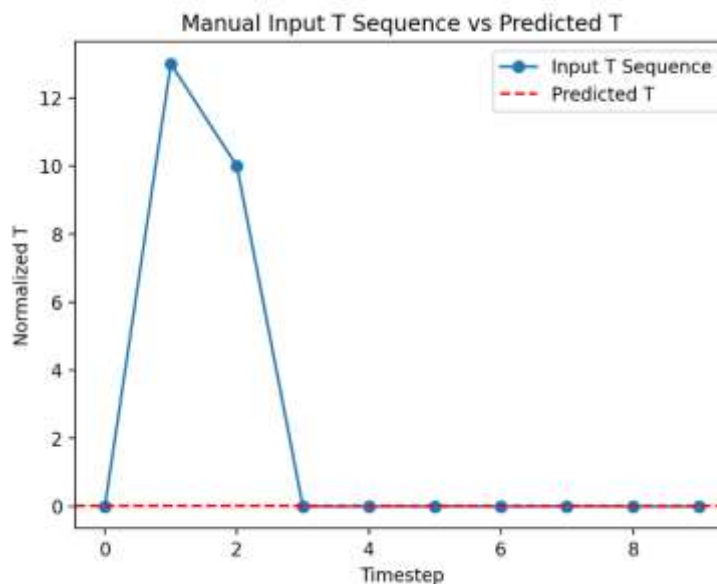


Fig: prediction graph

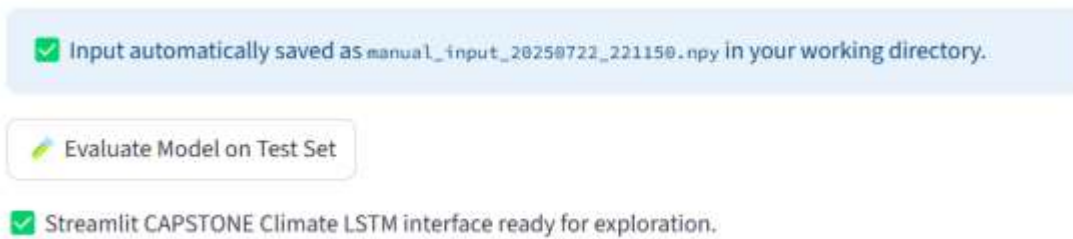


Fig: prediction

4. RESULTS:

The result of the project shows that the Physics-Informed BiLSTM model performed very well in predicting atmospheric variables like temperature (T) and geopotential height (Z). After 30 epochs of training, the model achieved high accuracy with low errors—such as a Mean Squared Error (MSE) of 0.0603 for temperature and 0.0125 for geopotential height, along with R^2 scores of 0.9391 and 0.9874 respectively. These results indicate that the model could accurately capture the temporal patterns in the data while remaining consistent with physical laws. The inclusion of physics-based loss improved the model's generalization and stability, especially in areas where traditional models often fail. Overall, the model demonstrated strong predictive performance and reliability for climate forecasting tasks.

5. DISCUSSION:

The project discusses how combining Physics-Informed Neural Networks (PINNs) with BiLSTM improves weather prediction by ensuring both accuracy and physical consistency. The use of physics-based loss alongside traditional loss helps reduce overfitting and improves generalization. Visual tools like heatmaps and scatter plots confirmed the model's reliability. Overall, the discussion highlights that integrating domain knowledge with machine learning leads to more stable and interpretable climate forecasting systems.

6. CONCLUSION

This project evaluated the effectiveness of a physics-informed BiLSTM model in predicting two key climate variables—temperature (T) and geopotential height (Z). The dataset was properly preprocessed and standardized to ensure fair evaluation using metrics like MSE, MAE, RMSE, and R^2 score, along with visual comparisons. The model showed excellent predictive performance, achieving a high R^2 score of 0.9874 for Z and 0.9391 for T, indicating strong accuracy and generalization. By including physics-based loss, the model followed natural laws, remained stable during training, avoided overfitting, and performed well even with limited data. Overall, the physics-informed BiLSTM successfully learned temporal patterns and produced accurate, physically consistent climate predictions.

7. FUTURE SCOPE:

The future scope of this project includes several areas for improvement and expansion. The model can be enhanced by adding more physical constraints and domain-specific rules to improve accuracy, especially when data is limited. Incorporating diverse datasets—such as simulated and experimental measurements—can make the model more robust under different conditions. The inclusion of external factors or boundary conditions could further improve prediction reliability. Enhancing the scalability and computational efficiency of the BiLSTM, possibly through lighter model architectures, would make it suitable for real-time applications. Additionally, testing the model in other domains like engineering systems or bioinformatics and integrating techniques like reinforcement learning and hybrid modeling could enable adaptive and flexible forecasting in dynamic environments. These improvements would broaden the usefulness of physics-informed machine learning for solving complex real-world problems.

8. ACKNOWLEDGEMENT:



Miss P. Bindhu priya working as an Assistant professor in Department of MCA in sanketika vidya parishad engineering college, Visakhapatnam, AP with 2.5 yrs teaching experience and member of IAENG, accredited by Naac with her areas of interests in C, Data warehousing and data mining, Design and analysis of algorithm, python, software engineering.



Mandavakuriti Harika is pursuing her final semester MCA in Sanketika Vidya Parishad Engineering College, accredited with A grade by NAAC, affiliated by Andhra University and approved by AICTE. With interest in Machine learning Mandavakuriti Harika has taken up her PG project on IMPLEMENTING WEATHER PREDICTION USING PHYSICS INFORMED NEURAL NETWORKS (PINNS) and published the paper in connection to the project under the guidance of P. BINDHU PRIYA, Assistant Professor, SVPEC.

REFERENCES

- [1] Physics-informed neural networks for high-resolution weather reconstruction from sparse weather stations
<https://pmc.ncbi.nlm.nih.gov/articles/PMC11308984/>
- [2] Physics-informed machine learning: case studies for weather and climate modelling
<https://royalsocietypublishing.org/doi/full/10.1098/rsta.2020.0093>
- [3] ClimODE: Climate and Weather Forecasting with Physics-informed Neural ODEs
<https://arxiv.org/abs/2404.10024>
- [4] Forecasting Buoy Observations Using Physics-Informed Neural Networks
<https://ieeexplore.ieee.org/abstract/document/10530311>
- [5] Accurate photovoltaic power prediction via temperature correction with physics-informed neural networks
<https://www.sciencedirect.com/science/article/abs/pii/S0360544225021887>
- [6] Ensemble Learning with Physics-Informed Neural Networks for Harsh Time Series Analysis
https://link.springer.com/chapter/10.1007/978-3-031-61418-7_5
- [7] A physics-informed neural network for turbulent wake simulations behind wind turbines
<https://pubs.aip.org/aip/pof/article-abstract/37/1/015110/3329123/A-physics-informed-neural-network-for-turbulent?redirectedFrom=fulltext>
- [8] Physics-Informed Neural Networks (PINN) for temporal forecasting of monsoon rainfall variability in southern peninsular India
<https://www.researchsquare.com/article/rs-6185189/v1>

- [9] M-ENIAC: A Physics-Informed Machine Learning Recreation of the First Successful Numerical Weather Forecasts
<https://agupubs.onlinelibrary.wiley.com/doi/full/10.1029/2023GL107718>
- [10] A numerical approach for soil microbiota growth prediction through physics-informed neural network
<https://www.sciencedirect.com/science/article/abs/pii/S0168927424002290>
- [11] PINT: Physics-Informed Neural Time Series Models with Applications to Long-term Inference on WeatherBench 2m-Temperature Data
<https://arxiv.org/abs/2502.04018>
- [12] Application of Physics-Informed Neural Networks for Solving the Inverse Advection-Diffusion Problem to Localize Pollution Sources
<https://arxiv.org/abs/2503.18849>
- [13] Potential of Using Physics with Artificial Intelligence: Physics-Infused Machine Learning and Neural Networks for Spatiotemporal Analysis of Hydro-climatic Events
https://link.springer.com/chapter/10.1007/978-3-031-76189-8_3
- [14] Self-scalable Tanh (Stan): Faster Convergence and Better Generalization in Physics-informed Neural Networks
<https://arxiv.org/abs/2204.12589>
- [15] Modeling Siltation of River Channels Using the Physics-Informed Neural Networks Method and Numerical Simulation
<https://www.espublisher.com/journals/article/details/1296>
- [16] An Expert's Guide to Training Physics-informed Neural Networks
<https://arxiv.org/abs/2308.08468>
- [17] An AI framework integrating physics-informed neural network with predictive control for energy-efficient food production in the built environment
<https://www.sciencedirect.com/science/article/abs/pii/S0306261923008140>
- [18] AQ-PINNs: Attention-Enhanced Quantum Physics-Informed Neural Networks for Carbon-Efficient Climate Modeling
<https://arxiv.org/abs/2409.01626>
- [19] PINN-Infused Hybrid ML Forecasting on Lake-Effect Precipitation
<https://www.inno-press.com/index.php/JAAI/article/view/26>
- [20] Physics-informed neural networks for solving incompressible Navier–Stokes equations in wind engineering
<https://pubs.aip.org/aip/pof/article-abstract/36/12/121303/3325454/Physics-informed-neural-networks-for-solving>