BILINGUAL PUBLISHING GROUP

Journal of Atmospheric Science Research

https://journals.bilpubgroup.com/index.php/jasr

ARTICLE

Hybrid Deep Learning and Bayesian Framework for Long-Term Fog Forecasting: A Case Study at Lucknow Airport in the Indo-Gangetic Plains

Deep Chaulya

Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna 801106, India

ABSTRACT

Fog significantly affects aviation, transportation, agriculture, and energy in North India during winter. Accurate prediction is difficult due to nonlinear meteorological interactions and uncertainty. Traditional models provide point estimates without uncertainty, while Bayesian methods offer probabilistic outputs but struggle to capture complex temporal dependencies in atmospheric time-series data effectively. To address these limitations, this study proposes a hybrid long-term forecasting framework that integrates a three-layer Bidirectional Long Short-Term Memory (BiLSTM) network with a Modified Bayesian Beta Regression model. The model is trained and evaluated using 24 years (2000–2023) of meteorological observations from Lucknow Airport, representing a fog-prone region within the Indo-Gangetic Plains. Fog behaviour was represented through a normalized Fog Index (0–1) that incorporates both intensity and persistence, offering a more stable and informative forecast target than raw visibility. The BiLSTM component captures temporal dependencies and produces accurate point forecasts, achieving a root mean square error (RMSE) of 0.125 for a three-day prediction horizon. The Bayesian layer enhances reliability by generating calibrated uncertainty intervals. The resulting model achieved an average interval width of approximately 0.59 and a prediction interval coverage close to 80%, effectively representing both aleatoric and epistemic uncertainty. Additional techniques, including Min-Max normalization, sequence windowing, Fast Fourier Transform (FFT) feature augmentation, learning-rate scheduling, and Monte Carlo dropout,

*CORRESPONDING AUTHOR:
Deep Chaulya, Department of Computer Science and Engineering, Indian Institute of Technology Patna, Patna 801106, India;
Email: deep_2408res08@iitp.ac.in

ARTICLE INFO
Received: 9 February 2026 | Revised: 9 March 2026 | Accepted: 23 April 2026 | Published Online: 28 April 2026
DOI: https://doi.org/10.30564/jasr.v9i2.13144

CITATION
Chaulya, D., 2026. Hybrid Deep Learning and Bayesian Framework for Long-Term Fog Forecasting: A Case Study at Lucknow Airport in the Indo-Gangetic Plains. Journal of Atmospheric Science Research. 9(2): 26–63. DOI: https://doi.org/10.30564/jasr.v9i2.13144

improved generalization and model stability. The proposed hybrid framework outperforms standalone models and shows strong potential for operational fog forecasting, offering both accurate predictions and uncertainty-aware confidence estimates for aviation, transportation, and early-warning systems.

Keywords: Fog Index Forecasting; BiLSTM; Deep Learning; Modified Bayesian Regression; Time-Series Analysis

1. Introduction

Fog is an atmospheric boundary-layer phenomenon formed due to the suspension of microscopic water droplets or ice crystals near the ground surface, resulting in a significant reduction in horizontal visibility, often below 1 km, and in severe cases below 200 m [1]. Although fog is a naturally occurring meteorological process, its impact on society, economy, and transportation is profound. In regions such as the Indo-Gangetic Plains (IGP) of India, fog formation during winter months, particularly between November and February is frequent and severe. Dense fog events in the IGP have repeatedly caused flight cancellations, aviation delays, railway traffic disruption, vehicular accidents, agricultural scheduling inefficiencies, and large-scale logistical failures, with measurable economic and human impact [2–5]. Due to its recurring intensity and persistence, fog has become one of the major environmental hazards affecting daily life, public safety, and regional economic productivity in Northern India.

1.1. Fog Climatology and Geographical Significance

Fog exhibits a strong geographical dependency, and its occurrence is closely linked to regional topography, surface characteristics, and atmospheric boundary-layer dynamics. Among the fog-affected regions globally, the IGP stands out as one of the most severely impacted fog belts during winter. Extending longitudinally across Punjab, Haryana, Delhi, Uttar Pradesh, Bihar, and northern Rajasthan of India, the region forms a vast low-lying alluvial basin bordered by the Himalayas to the north and the Thar Desert to the west. Due to this unique geographical configuration, cold continental air masses originating from the Himalayas interact with moist low-level air from the IGP, resulting in ideal atmospheric conditions for fog formation and persistence [3].

The fog climatology of the IGP is highly characteristic and distinct from other global fog-prone regions such as Eastern China, California's Central Valley, or European coastal plains. A combination of meteorological, synoptic, and anthropogenic variables regulates fog behaviour in this region. The frequent formation of fog in the region is primarily driven by a combination of meteorological and anthropogenic factors. Low nocturnal temperatures resulting from strong radiative cooling during winter nights frequently bring the near-surface air temperature close to the dew point, creating conditions favourable for condensation. High relative humidity and a narrow dew point spread further enhance the likelihood of fog formation by increasing the saturation potential of the air. Additionally, the presence of a stable surface boundary layer and strong temperature inversion traps moisture and pollutants near the ground, preventing vertical mixing. Weak wind circulation and atmospheric stagnation contribute to the persistence of fog by inhibiting dispersal of suspended droplets. Finally, elevated concentrations of aerosols and particulate matter originating from industrial emissions, vehicular exhaust, and biomass burning act as effective cloud condensation nuclei, intensifying droplet formation and increasing both the density and duration of fog events. These interacting variables collectively support the development of fog events lasting from 6 to 72 h, forming persistent episodes often referred to as cold wave fog spells. Such long-duration fog events are especially prevalent from late December to mid-January, aligning with peak winter inversion and minimum annual temperatures.

In addition to spatial variability, fog in the IGP demonstrates strong inter-annual and intra-seasonal variability, influenced by large-scale atmospheric phenomena such as the El Niño–Southern Oscillation (ENSO), Arctic Oscillation, and regional aerosol load variations. Long-term observations collected from the India Meteorological Department (IMD), surface synoptic stations, and airport meteorological logs reveal significant year-to-year fluctuations in fog intensity, duration, and timing. This variability highlights the dynamic nature of fog formation processes in the region and the influence of evolving climatic drivers.

Figure 1 illustrates the long-term trend of winter fog days recorded in Delhi between 2000 and 2025. The data show alternating cycles of increasing and decreasing fog frequencies, demonstrating how factors such as urbanization, aerosol concentration, moisture availability, and boundary-layer meteorology affect fog climatology. For instance, phases of rapid industrialization and biomass burning have been associated with increased fog occurrences, whereas warmer winters and improved fuel standards have correlated with reduced fog days in some years.

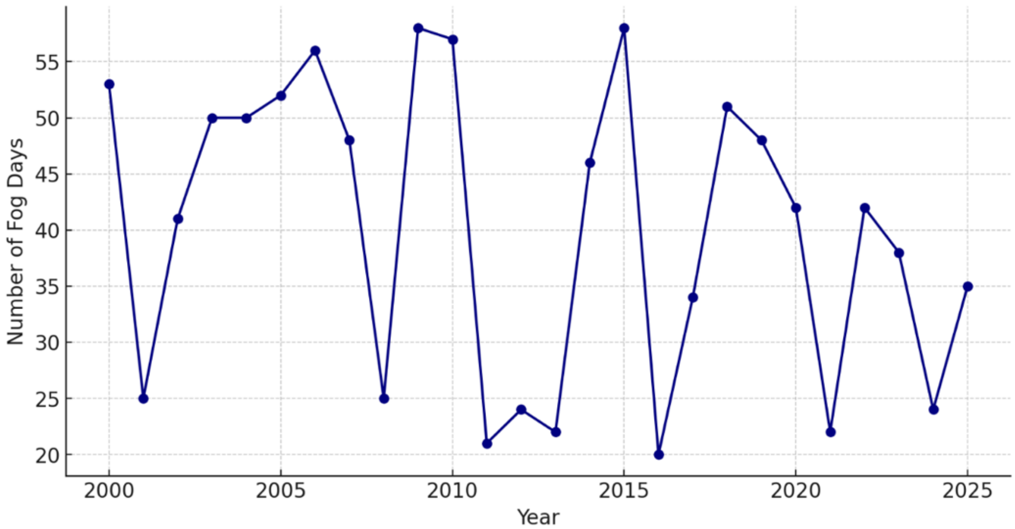

Figure 1. Inter-annual variation in winter fog days over Delhi from 2000 to 2025, illustrating long-term trends, seasonal persistence, and year-to-year fluctuations in fog frequency associated with changing boundary-layer conditions and regional atmospheric dynamics.

Such climatological complexity reinforces the need for advanced fog forecasting methodologies that account not only for historical meteorological trends but also for evolving atmospheric dynamics and anthropogenic influences. Understanding the geographical significance and spatial–temporal behaviour of fog within the IGP is therefore essential in designing predictive frameworks capable of supporting public safety, transportation scheduling, disaster preparedness, and environmental policy planning.

1.2. Socioeconomic Impact of Fog

Dense fog exerts a profound influence on socioeconomic systems, primarily due to its ability to significantly reduce visibility, often to critical operational thresholds below 200 m [6], and in extreme scenarios to less than 50 m, thereby compromising public mobility and safety [7]. As visibility decreases and environmental uncertainty increases, human decision-making, transport control systems, and automated navigation mechanisms face elevated operational risks. **Figure 2** presents a conceptual representation of how declining visibility correlates with increasing disruption severity across transportation and daily life activities, demonstrating that even moderate fog conditions can lead to measurable delays, while dense fog conditions may result in complete operational shutdowns. Among all affected systems, aviation experiences the most immediate and widespread disruptions. Major airports across northern India, include Indira Gandhi International Airport (Delhi), Chandigarh Airport, Lucknow Airport, and Amritsar Airport. It recorded frequent runway closures, flight rescheduling, diverted landings, and in some cases, mass cancellations due to dense fog events. The consequences extend beyond logistical inconvenience: passenger safety protocols, fuel consumption, air traffic management, crew duty cycles, and airport resource allocation are all affected. During peak winter months, fog may ground aircraft for hours, affecting thousands of passengers and leading to significant economic losses for airlines and aviation authorities [8,9].

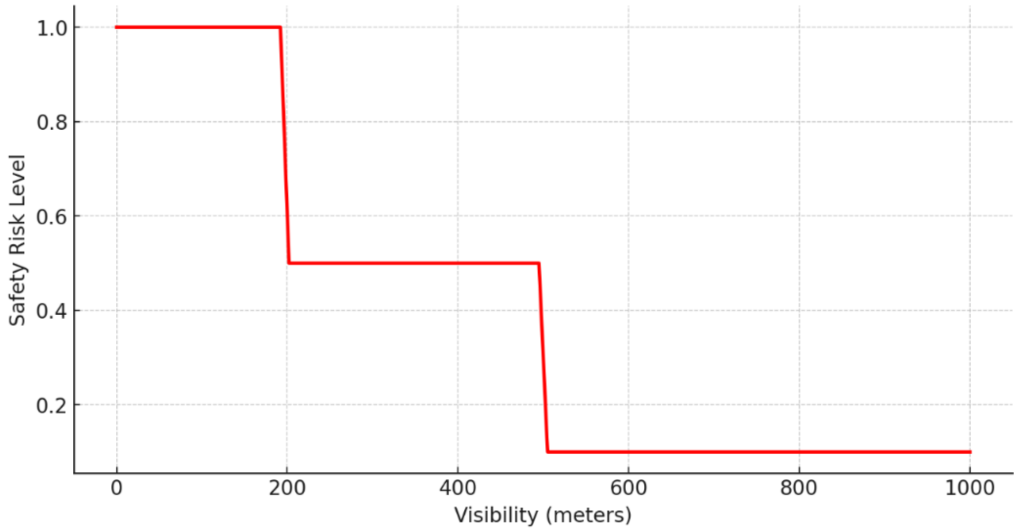

Figure 2. Conceptual relationship between fog intensity and transportation safety risk, illustrating how declining visibility levels correspond to increasing operational disruption and hazard severity across transportation systems.

Fog also represents a major hazard for road transportation, where abrupt reductions in visibility increase the likelihood of accidents, multi-vehicle collisions, and traffic congestion. Empirical records from northern India indicate a sharp rise in fog-related accidents during December and January, particularly along high-speed corridors such as National Highways (NH-24, NH-44, NH-19) and expressways including the Yamuna Expressway and Agra–Lucknow Expressway [10]. Reduced lane visibility, slower driver reaction time, impaired depth perception, and unpredictable traffic behaviour amplify the risk, often leading to chain-reaction crashes involving multiple vehicles. This not only endangers human life but also places an additional burden on emergency response, medical services, and insurance systems.

Similarly, railway operations, which serve as the backbone of long-distance mobility in the IGP, are severely affected by fog events. Indian Railways reports frequent delays during winter fog, especially along major northern rail routes such as Delhi-Howrah and Delhi-Amritsar corridors. Trains operating under dense fog conditions are forced to reduce speed due to compromised visibility of signals and track infrastructure, resulting in cascading delays that affect scheduling on a national scale. Long stoppages, rerouting, station congestion, and disrupted freight movement further intensify systemic inefficiencies and contribute to financial losses. Beyond transportation, dense fog has additional indirect socioeconomic consequences. It affects logistics and supply chains, delaying freight delivery and disrupting perishables, pharmaceutical logistics, and agricultural distribution networks. In the agricultural sector, persistent fog reduces incoming solar radiation, lowering crop temperatures, slowing photosynthesis, and increasing susceptibility to fungal infections in crops such as wheat, mustard, potato, and sugarcane [11]. Prolonged cold fog periods may also delay harvesting and irrigation cycles, affecting overall productivity and revenue.

Fog impacts also extend to public health and daily life. Reduced sunlight combined with elevated particulate concentrations intensifies respiratory disease incidence, particularly among vulnerable populations. Schools, factories, and outdoor workforce operations may experience reduced working hours due to prolonged foggy mornings. In addition, fog-induced disruptions impose psychological stress, commuting uncertainty, and reduced operational confidence among the population.

1.3. Challenges in Fog Prediction

Fog prediction remains one of the most scientifically and operationally challenging tasks in meteorology due to the complex interplay of microphysical, thermodynamic, and atmospheric boundary-layer processes. Unlike large-scale weather systems such as cyclones or monsoons, fog forms within a shallow layer close to the surface often only 10 to 300 m above ground, where atmospheric behaviour is highly sensitive to minute changes in temperature, humidity, and radiation balance. This micro-scale nature introduces substantial uncertainty in modelling and predicting fog onset, duration, and dissipation. One of the foremost difficulties arises from small-scale vertical stratification, where subtle variations in temperature gradients and moisture concentration dictate whether atmospheric air parcels reach saturation. Even small measurement errors in temperature or dew point can shift predictions from fog formation to clear-air conditions.

Fog formation is further influenced by boundary-layer microphysics, including condensation processes, radiative cooling, droplet activation, and turbulence suppression. These processes interact nonlinearly with surface energy balance and atmospheric stability, making them difficult to parameterize accurately in forecasting systems. The temporal evolution of fog adds another layer of complexity. Fog can form rapidly within minutes and dissipate even more abruptly once turbulence or radiative heating is introduced. This rapid transition between stages of fog evolution requires forecasting systems capable of high temporal resolution and accurate nowcasting capabilities, which many traditional models lack.

Additionally, fog behaviour is strongly modulated by aerosols, synoptic forcing, and topography. Aerosols serve as condensation nuclei and influence droplet size spectra, affecting fog density and persistence. In heavily industrialized or polluted regions such as the IGP, elevated concentrations of particulates significantly amplify fog stability and duration. Synoptic-scale meteorological systems, including Western Disturbances, surface inversions, and moisture advection events, further complicate fog prediction by altering thermodynamic equilibrium conditions. Topographic settings such as valleys or areas near rivers also trap cold air and moisture, increasing fog likelihood. Collectively, these dynamic and interdependent factors contribute to both aleatoric uncertainty (natural atmospheric variability) and epistemic uncertainty (limitations in model structure or measurement accuracy), making fog prediction highly uncertain. Because of these difficulties, predictive approaches for fog have evolved across three methodological domains: physics-based forecasting, statistical modelling, and machine learning approaches. Each method offers unique advantages but also exhibits limitations that restrict its standalone reliability in operational forecasting environments. **Table 1** summarizes the comparative limitations and capabilities of these methods.

Table 1. Comparative summary of major fog forecasting approaches, outlining commonly used methods along with their strengths and key limitations in predicting fog dynamics.

Category	Common Methods Used	Strengths	Key Limitations for Fog Forecasting
Physics-Based Models	Numerical Weather Prediction (NWP), Weather Research and Forecasting (WRF), mesoscale boundary-layer models	Physically interpretable; simulates thermodynamics and microphysics; integrates large-scale weather systems	Computationally expensive; requires high-resolution land-surface and aerosol inputs; parameterization uncertainty; coarse grid often fails to capture shallow fog layers
Statistical Models	Logistic regression, ARIMA, discriminant analysis, persistence models	Computationally fast; suitable for operational use; requires less input data	Assumes linearity and stationarity; oversimplifies nonlinear fog dynamics; low generalization ability; sensitive to noise
Machine Learning Models	Support Vector Machines (SVM), Random Forests (RF), Gradient Boosting (XGBoost, LightGBM)	Handles nonlinear relationships; strong classification performance; less reliance on explicit atmospheric equations	Treats samples independently; limited ability to learn sequential or temporal dependencies; performance degrades for long-term forecasting

Although numerical weather prediction models produce detailed physical simulations, their performance in visibility prediction is commonly constrained by coarse spatial resolution and uncertainty in representing land-surface exchange processes, radiative cooling, and droplet microphysics. The computational demand further restricts their suitability for high-frequency real-time operational forecasting. On the other hand, statistical models provide fast predictions and require fewer computational resources; however, they rely on assumptions that fail under nonlinear and non-stationary atmospheric behaviour. Machine learning models represent a more advanced alternative capable of learning nonlinear correlations among meteorological predictors, yet they generally lack the intrinsic ability to model sequence-dependent time-series behaviour, which is critical in fog evolution where past conditions strongly influence future states.

1.4. Evolution towards Deep Learning-Based Fog Forecasting

Fog forecasting has evolved significantly over the past two decades, transitioning from empirical rule-based forecasting to physics-driven numerical simulations and, more recently, to data-driven computational intelligence methods. This shift has been largely enabled by the growing availability of high-resolution meteorological observations, remote sensing data, reanalysis products, and automated weather station (AWS) archives. The increasing volume, variety, and temporal resolution of atmospheric datasets have provided the foundation for machine learning and deep learning algorithms to extract hidden patterns and model fog behaviour with improved accuracy and generalizability. Early data-driven attempts in fog prediction were dominated by classical machine learning approaches such as Support Vector Machines (SVM), Random Forest (RF), Gradient Boosting, and Logistic Regression, which demonstrated notable improvements over deterministic and statistical visibility forecasting methods, particularly for fog/no-fog classification tasks in high-risk regions such as the IGP [11]. While these models successfully captured nonlinear relationships among atmospheric variables, their reliance on feature engineering, manual input selection, and independence assumptions between samples limited their capability in long-term fog intensity prediction. Most classical machine learning models treat each observation as a discrete point, thereby neglecting temporal continuity, memory structure, and the sequential evolution of meteorological conditions components essential to modelling fog lifecycle dynamics.

The introduction of deep learning architectures marked a transformative milestone in fog forecasting research. Deep neural networks, unlike traditional machine learning algorithms, are capable of automated feature representation and hierarchical learning. Among these architectures, Recurrent Neural Networks (RNNs) were among the earliest to be adopted for weather-related sequence modelling due to their ability to process time-dependent input. However, basic RNNs suffer from gradient degradation problems that restrict their ability to learn long-range dependencies within meteorological time series. To address these limitations, Long Short-Term Memory (LSTM) networks were introduced, allowing models to retain useful temporal information over longer intervals, mitigating gradient vanishing issues, and enabling robust pattern learning from extended atmospheric sequences. LSTM architectures became particularly valuable for fog forecasting, as fog formation is influenced by gradual changes in meteorological variables such as dew point depression, relative humidity build-up, radiative cooling, and inversion layer stability, all of which evolve over multiple hours rather than instantaneously. Studies incorporating LSTM networks reported significant improvements in fog onset prediction accuracy and temporal consistency compared to traditional machine learning and statistical approaches.

Further advancements were made with the development of Bidirectional LSTM (BiLSTM) architectures, which enhanced performance by processing meteorological sequences in both forward and backward temporal directions. This bidirectional learning capability allows BiLSTM models to capture delayed and anticipatory dependencies, providing a more comprehensive representation of atmospheric evolution patterns compared to standard unidirectional models [12]. In fog prediction scenarios, where pre-formation environmental cues and dissipation signals both influence forecast confidence BiLSTM models have demonstrated superior skill in modelling uncertainty and capturing transitional phases between different fog intensity levels. Progress in spatial-temporal deep learning further expanded the capability of predictive modelling. Hybrid architectures such as Convolutional Neural Network-LSTM (CNN-LSTM) frameworks combine the spatial feature extraction ability of CNNs with the sequence modelling strengths of LSTM networks, making them particularly effective for visibility and fog prediction using gridded meteorological fields, reanalysis datasets, satellite imagery, and multisensor fusion inputs [13]. These hybrid models are capable of jointly learning topographic context, boundary-layer structure, and temporal evolution without requiring extensive manual predictor design, marking a major advancement toward automated, scalable fog forecasting pipelines.

More recent research directions have focused on probabilistic deep learning frameworks, uncertainty-aware models, and hybrid architectures integrating Bayesian inference, transformer networks, and ensemble learning to address aleatoric and epistemic uncertainty. Such developments reflect a paradigm shift from simple visibility classification toward robust multi-step, confidence-aware fog intensity forecasting, suitable for real-world operational deployment in transportation safety, aviation management, and early warning systems.

1.5. Hybrid Deep Learning and Bayesian Framework

While deep learning-based models have demonstrated significant advancements in fog forecasting accuracy, a major limitation persists: most deep neural networks generate deterministic point forecasts without providing any measure of prediction confidence. In safety-critical sectors, such as aviation, intelligent transportation systems, railway scheduling, and automated vehicular navigation forecast reliability is just as important as forecast accuracy. A single predicted fog severity value offers limited operational usefulness unless decision makers also understand the degree of uncertainty associated with that prediction. For instance, forecasting a fog index of 0.70 may require drastically different response strategies depending on whether the associated uncertainty range is narrow (0.68–0.72) or wide (0.45–0.94). Thus, credible interval estimation and uncertainty awareness are essential components of any robust fog forecasting framework.

Bayesian learning methods naturally address this need by producing probabilistic outputs rather than fixed deterministic values. Bayesian models estimate posterior distributions over model parameters, thereby quantifying both aleatoric uncertainty (inherent randomness in fog-generating atmospheric processes) and epistemic uncertainty (uncertainty arising from limited or imperfect training data). Through these probabilistic representations, Bayesian regression models can generate prediction intervals, confidence bounds, and likelihood-based risk metrics features that are particularly valuable in weather forecasting applications where uncertainty must be transparently communicated to operational stakeholders. However, despite these strengths, standalone Bayesian regression techniques are not well-suited to modelling complex nonlinear atmospheric evolution. Their limited capacity for capturing long-term temporal patterns, memory structures, delayed meteorological feedback loops, and multivariate sequential dependencies restricts their effectiveness for multi-step fog time-series forecasting. To reconcile the strengths and limitations of both paradigms, this study proposes a hybrid fog forecasting architecture that integrates deep temporal modelling with Bayesian uncertainty quantification. The hybrid framework employs a three-layer BiLSTM network as the core sequence-learning engine, enabling the system to capture nonlinear dependencies across multiple time steps while learning both historical and emerging atmospheric patterns. The bidirectional design allows the model to interpret meteorological evolution in both forward and backward directions, producing richer feature embedding that reflects the full temporal context of fog development.

To complement the BiLSTM's deterministic forecasting capability, the model incorporates a Modified Bayesian Beta Regression layer, specifically optimized for forecasting the Fog Index, a variable constrained between 0 and 1. This Bayesian component generates calibrated probabilistic outputs and credible intervals, ensuring the system provides uncertainty-aware forecasts suitable for operational risk evaluation. Model calibration techniques, including beta-likelihood optimization and posterior distribution shaping, ensure reliable uncertainty representation across multiple lead times. Beyond the core architecture, several auxiliary mechanisms are incorporated to enhance the stability and generalization ability of the system. Monte Carlo (MC) dropout is applied during inference to simulate model variability and approximate Bayesian behaviour within the deep learning component, thereby strengthening epistemic uncertainty estimation. Learning-rate scheduling dynamically adjusts model optimization behaviour to avoid overfitting and ensure convergence stability during long training cycles. Additionally, Fast Fourier Transform (FFT)-based feature augmentation is applied to extract spectral signatures and periodic behaviours embedded in historical meteorological signals, enabling the model to learn seasonal patterns and recurrent winter fog cycles more effectively.

In this study, the modification in Bayesian Beta Regression is not limited to a single element (e.g., prior choice or MC Dropout alone), but instead represents a structural reformulation of standard Bayesian Beta regression to make it suitable for sequential fog forecasting. Specifically, the proposed modification consists of three integrated changes to classical Bayesian Beta regression:

(a) Latent Temporal Conditioning via BiLSTM: In standard Bayesian Beta regression, the response variable y_t is modeled as:

$$y_t \sim Beta\left(\alpha\left(X_t\right), \beta\left(X_t\right)\right) \tag{1}$$

where the shape parameters depend on static predictors X_t.

In the proposed model, the Beta parameters are conditioned on the latent temporal state learned from meteorological sequences:

$$h_t = BiLSTM\left(X_{t-L:t}\right) \tag{2}$$

$$y_t \sim Beta\left(\alpha\left(h_t\right), \beta\left(h_t\right)\right) \tag{3}$$

Thus, the regression layer is no longer linked to raw inputs, but to a sequence-informed hidden representation.

(b) Nonlinear Parameter Mapping: Instead of linear link functions traditionally used in Beta regression, the shape parameters are obtained via nonlinear transformations:

$$\alpha_t = f_\alpha\left(h_t\right), \quad \beta_t = f_\beta\left(h_t\right) \tag{4}$$

ensuring positivity through softplus activation.

(c) Approximate Bayesian Inference via MC Dropout: Rather than relying on computationally expensive MCMC sampling, Monte Carlo Dropout has been adapted to approximate posterior predictive uncertainty:

$$\widehat{y_t} = \frac{1}{N}\sum_{i=1}^{N} f\left(h_t; \theta_i\right) \tag{5}$$

where θ_i represents dropout-perturbed weights.

This combination transforms the regression component into a temporal probabilistic decoder, which justifies the use of the term modified.

Conceptually, the modification replaces the static regression assumption of standard Bayesian Beta regression with a sequence-aware probabilistic framework.

Standard formulation:

$$y_t \sim Beta\left(\mu_t\phi, \left(1-\mu_t\right)\phi\right) \tag{6}$$

$$\mu_t = g^{-1}\left(X_t\beta\right) \tag{7}$$

Proposed formulation:

$$h_t = BiLSTM\left(X_{t-L:t}\right) \tag{8}$$

$$\mu_t = \sigma\left(W_\mu h_t\right) \tag{9}$$

$$\phi_t = softplus\left(W_\phi h_t\right) \tag{10}$$

$$y_t \sim Beta\left(\mu_t\phi_t, \left(1-\mu_t\right)\phi_t\right) \tag{11}$$

Predictive uncertainty is then estimated through MC Dropout:

$$p\left(y_t|X\right) \approx \frac{1}{N}\sum_{i=1}^{N} Beta\left(\alpha_i, \beta_i\right) \tag{12}$$

The 'Modified' Bayesian Beta Regression differs from standard formulations in three ways:

(a) Latent Temporal Conditioning: The shape parameters α and β are conditioned on the BiLSTM hidden state h_t rather than raw inputs X_t.

(b) Nonlinear Parameter Mapping: It employs a Softplus activation to ensure shape parameters remain strictly positive and capture nonlinear saturation.

(c) Approximate Inference: It replaces expensive MCMC sampling with Monte Carlo (MC) Dropout to estimate the posterior predictive distribution $p\left(y_t|X\right)$ efficiently for operational use.

1.6. Forecast Variable: Fog Index

Selecting an appropriate forecast variable is a critical component of fog prediction modelling, as it directly influences model interpretability, temporal stability, and predictive performance. Traditionally, most operational fog forecasting systems rely on raw visibility measurements reported in meters or kilometers. Although visibility is an intuitive and operationally useful indicator, it is highly volatile and exhibits abrupt fluctuations in response to micrometeorological changes such as momentary turbulence, minor radiative balance shifts, or transient aerosol concentration changes. Such fluctuations introduce noise into the modelling process, increasing prediction uncertainty and reducing long-term forecast stability. Moreover, raw visibility values often fail to reflect the continuous nature of fog evolution, especially when fog persists for extended durations or exhibits multi-phase development, including pre-conditioning, dense formation, maintenance, and gradual dissipation phases.

To address these limitations, the present study adopts a Fog Index, a normalized metric ranging between 0 and 1 designed to represent fog behaviour more holistically. In contrast to visibility by itself, the Fog Index combines both the intensity of fog, as determined by the reduction in visibility, and the duration of fog, which indicates how long the fog conditions last over a period of time. This combined representation captures the temporal continuity and severity scale of fog events, making it especially suitable for modelling extended lead times where atmospheric memory and gradual transitions play a significant role. By compressing raw visibility into a bounded scale, the Fog Index mitigates issues associated with extreme values, non-stationarity, and rapid volatility, thereby enabling smoother learning dynamics for deep learning architectures. The Fog Index also offers distinct computational benefits. Its normalized structure enhances the performance of machine learning and deep learning algorithms by reducing numerical variability and improving convergence behaviour during training. Models such as LSTM and BiLSTM rely heavily on sequential pattern recognition; thus, a stable and continuous target variable like the Fog Index preserves the underlying temporal structure more effectively than discrete or highly fluctuating visibility measurements. Furthermore, the Fog Index serves as an appropriate response variable for Bayesian Beta Regression, given that the beta probability distribution is defined on the (0,1) interval and is commonly used for modelling bounded environmental indices, proportions, and normalized physical phenomena.

Another important advantage of this variable lies in its suitability for forecasting multi-day lead times, particularly the 1–5 day prediction horizon considered in this study. Long-term fog prediction requires that models capture seasonal patterns, gradual boundary-layer cooling, synoptic variability, and cumulative atmospheric moisture buildup. The Fog Index, by emphasizing persistence and intensity trends, offers more predictable temporal trajectories than raw visibility, thereby improving regression stability and enabling more reliable medium-range forecasts.

1.7. Contributions of This Study

Fog remains a highly challenging atmospheric phenomenon to forecast due to its localized scale, nonlinear formation dynamics, and strong dependence on boundary-layer processes. Although substantial progress has been made using physical models, machine learning, and hybrid frameworks, gaps continue to exist in long-term fog prediction reliability, uncertainty quantification, and operational applicability. In response to these challenges, the present research presents a set of methodological and scientific contributions aimed at strengthening both the forecasting capability and the interpretability of fog modelling systems. The major contributions of this study are summarized below:

First, the study provides a comprehensive analysis of the evolution of fog forecasting methodologies, spanning classical statistical techniques, standalone machine learning models, and recent advances in deep learning. This review consolidates findings across decades of research and identifies performance gaps in traditional forecasting pipelines, particularly in relation to temporal dependency handling, uncertainty modelling, and long-horizon prediction accuracy. By synthesizing this evolving research landscape, the study offers a structured reference for future fog prediction research and operational system development.

Second, this research introduces a novel hybrid fog forecasting framework that integrates a three-layer BiLSTM deep learning network with a modified Bayesian Beta Regression model. The BiLSTM component enables robust learning of nonlinear and sequential atmospheric dependencies, while the Bayesian extension addresses a critical limitation of existing deep learning models: the absence of uncertainty awareness. Through this hybridization, the framework generates both deterministic fog index predictions and calibrated probabilistic confidence intervals, making the system suitable for decision-support environments where risk evaluation is essential.

Third, the proposed model is rigorously tested using an extensive 24-year meteorological dataset (2000–2023) collected from Lucknow Airport, a location situated in the fog-prone IGP of India. Although the IGP is discussed for contextual relevance, the present study is based on data from a single station (Lucknow Airport). Therefore, results represent a localized case study rather than a region-wide generalization. The dataset includes visibility measurements, atmospheric variables, and derived Fog Index values. Long-term testing ensures that the evaluation captures seasonal variability, inter-annual climate fluctuations, aerosol influence shifts, and regional boundary-layer changes. Comparative experiments demonstrate the superior performance of the proposed hybrid model relative to baseline machine learning and deep learning architectures.

Fourth, the study performs a detailed analysis of feature contributions, seasonal variability patterns, spectral behaviour using Fast Fourier Transform (FFT), and model sensitivity characteristics. Such analysis provides deeper insight into the factors governing fog formation and supports model explainability an essential attribute when developing forecasting systems for high-risk applications.

Finally, the research offers practical recommendations for operational deployment, including strategies for real-time integration with ground stations, early warning dissemination pathways, computational optimization considerations, and potential applicability in intelligent transportation systems, automated decision support platforms, and aviation scheduling workflows. These recommendations highlight the feasibility of scaling the proposed hybrid model into a field-ready forecasting tool.

2. Literature Survey

Fog forecasting has been an active research domain for several decades due to its significant implications for aviation safety, transportation management, and public warning systems. Numerous scientific studies have explored fog dynamics using a wide spectrum of observational techniques, analytical models, and machine learning approaches. Early research in this field primarily relied on remote sensing datasets such as MODIS, INSAT, and ground-based visibility networks to detect fog occurrence and assess its spatial and temporal variability [14–18]. While these studies played a foundational role in mapping fog distribution through satellite observation [19] and identifying critical fog regimes, many remained focused on detection rather than multi-step forecasting.

A smaller subset of research incorporated Geographic Information Systems (GIS) to analyse fog behaviour, including dissipation processes, spatial patterns of fog intensity, and its seasonal recurrence [20,21]. Among the notable contributions, Saraf et al. [5] applied a combined remote sensing and GIS framework to identify highly fog-prone belts, investigate migration behaviour of fog layers, and generate short-term prediction maps for Northern India. Their work highlighted the IGP as a high-impact fog corridor and provided an important reference framework for subsequent modelling-based studies.

2.1. Deep Learning for Fog and Visibility Forecasting

In recent years, the rapid expansion of computational capabilities and access to high-resolution meteorological datasets has accelerated the adoption of deep learning approaches for fog and visibility forecasting. Unlike traditional statistical or numerical methods that rely heavily on explicit physical assumptions and parameterization, deep learning models are capable of learning complex nonlinear patterns embedded within atmospheric systems directly from data. This makes them particularly effective for fog prediction, as fog formation is influenced by subtle interactions among thermodynamic, radiative, boundary-layer, and aerosol-driven processes many of which are difficult to mathematically formalize in rule-based forecasting systems.

Early studies in this domain prominently utilized Convolutional Neural Networks (CNNs), especially in applications involving spatial datasets such as sky-camera videos, satellite imagery, or model-derived fog mask products. CNN-based approaches demonstrated strong capability in detecting fog signatures through the extraction of spatial texture patterns, stratiform cloud structures, thermal gradients, or low-level cloud base morphology observable in remote sensing imagery. Although CNNs have been successful for fog detection and classification, they inherently lack temporal memory and therefore struggle to forecast fog progression over extended time horizons, as the evolving atmospheric context cannot be fully represented through static spatial frames [22–24].

To address the temporal component of fog evolution, researchers increasingly turned toward sequential models, particularly Recurrent Neural Networks (RNNs) and their advanced gated variants such as LSTM and Gated Recurrent Unit (GRU) networks. These architectures are specifically designed to preserve temporal information over long input sequences, making them well-suited for forecasting fog using meteorological time-series variables including temperature, relative humidity, dew point depression, wind speed, pressure, and visibility. LSTM networks demonstrated notable improvements over classical regression and machine learning methods because of their ability to retain long-term atmospheric signals, such as the gradual accumulation of near-surface moisture or prolonged nocturnal radiative cooling leading to saturation [25,26]. Despite their forecasting strength, LSTM models often require larger datasets to avoid instability or overfitting, and challenges remain regarding interpretability when applied in high-stakes operational environments such as aviation forecasting.

As research progressed, the need to jointly model the spatial and temporal characteristics of fog led to the development of hybrid architectures such as Convolutional LSTM (ConvLSTM). These models inherit the spatial feature extraction capability of CNNs and combine it with the sequential memory of LSTMs, enabling simultaneous learning of evolving weather imagery sequences or gridded atmospheric fields. ConvLSTM models demonstrated superior forecasting skill compared to either CNN or LSTM architectures operating alone, particularly in applications where fog properties evolve dynamically over space and time, such as coastal marine fog, airport-scale low-visibility systems, and widespread radiation fog observed in Indo-Gangetic winter seasons [27,28]. The success of these architectures represents a significant step forward in operational fog forecasting systems.

In parallel, autoencoder-based architectures, including standard Autoencoders (AEs) and Variational Autoencoders (VAEs) have been increasingly applied to compress high-dimensional environmental data, reduce noise, extract latent atmospheric variables, and fuse multimodal observation sources such as satellite products, aerosol measurements, and station-level visibility observations [29,30]. Although autoencoders are rarely used as full standalone forecasting models, they play a critical role in preprocessing pipelines for downstream models, improving efficiency and reducing uncertainty by isolating high-value predictive features.

More recently, the emergence of transformer-based architectures and attention-driven deep learning has shifted the research landscape. Unlike recurrent models, transformers operate without sequential recurrence and instead use attention mechanisms to directly compute relevance relationships between alltime steps in the input sequence. This ability to model long-range dependencies without memory decay has positioned transformers as highly promising models for next-generation fog and weather forecasting. Furthermore, their interpretability achieved through attention weight visualization makes them attractive for aviation and transportation decision-support systems, where forecasting decisions must be explainable. However, these models tend to be computationally expensive and require extensive labelled datasets, making operational implementation still limited in fog-specific applications [31–33].

Overall, the evolution of deep learning methods has significantly advanced fog forecasting capability, transitioning from simple pattern recognition approaches toward sophisticated spatial-temporal reasoning frameworks capable of learning the atmospheric progression of fog formation, intensification, and dissipation. While each methodological family of CNNs, LSTMs, ConvLSTMs, AEs, and transformers, addresses fog prediction from a different computational perspective, emerging evidence suggests that hybrid and multimodal deep learning models will form the foundation of future operational fog prediction systems as more comprehensive meteorological datasets and computational resources become available.

2.2. Comparative Overview of Deep Learning Models

The evolution of fog forecasting research demonstrates a progressive shift from traditional computational methods to advanced deep learning architectures capable of modelling highly nonlinear atmospheric processes. As fog formation depends on a complex interplay of thermodynamics, moisture availability, aerosol interactions, radiative cooling, and synoptic-scale forcing, the choice of predictive architecture significantly impacts accuracy, interpretability, and operational applicability. A comparative evaluation of major deep learning families used in fog and visibility prediction is presented in **Table 2**, summarizing their input requirements, methodological strengths, and associated limitations.

Early deep learning applications in meteorological visibility prediction primarily adopted RNNs and their gated extensions LSTM and GRU architectures. These models were effective for sequential forecasting tasks because they could maintain temporal context, allowing them to detect slow atmospheric transitions leading to fog onset. For example, phenomena such as dew point convergence, boundary-layer cooling, and nocturnal inversion strength are gradual and recur across diurnal cycles, making LSTM-based modelling advantageous. However, despite their strong temporal learning capability, recurrent architectures face training challenges such as vanishing gradients and reduced accuracy for extended multi-day predictions, particularly when applied to sparse and seasonally variable datasets [34,35].

Table 2. Comparative overview of deep learning architectures used for fog and visibility forecasting, highlighting their input types, modelling strengths, and key limitations in capturing spatiotemporal fog dynamics.

Model Type	Typical Inputs	Key Strengths	Limitations/Remaining Gaps	Representative Studies
LSTM/GRU (RNNs)	Surface meteorological time series (temperature, humidity, dew point, wind, pressure, visibility), NWP variables	Effective in learning temporal dependencies; widely used for short-term fog forecasting	Limited spatial awareness; performance declines for long-horizon forecasts; often deterministic with limited uncertainty modelling	Miao et al. [34], Lu et al. [35]
CNN (1D/2D/3D)	Sky images, satellite imagery, CCTV frames, gridded NWP data	Strong spatial feature extraction; useful for visibility classification	Weak temporal memory; requires large labelled datasets; limited persistence modelling	Kamangir et al. [36], Chen et al. [37]
Temporal Convolutional Networks (TCN)	Time-series meteorological data	Efficient long-range sequence learning; parallelizable training	Limited ability to represent spatial diffusion of fog	Lu et al. [35]
Graph Neural Networks (GCN/ST-GNN)	Multi-station environmental networks	Captures spatial propagation of fog events	High model complexity; requires structured station networks; limited operational deployment	Qu et al. [38]
Transformers/Attention Models	Multimodal datasets (satellite + time-series)	Strong long-range temporal modelling; interpretable attention	Data-intensive; still emerging in fog forecasting applications	Qu et al. [38]
Hybrid CNN-RNN/ConvLSTM	Satellite image sequences + meteorological series	Improved spatiotemporal representation	Computationally demanding; stability challenges	Xiang et al. [27], Pagowski et al. [28], Kamangir et al. [36], Chen et al. [37]
Ensemble/Stacked Deep Models	Multi-source fusion (CNN + LSTM + GNN + NWP outputs)	High robustness and accuracy	Reduced interpretability; difficult operational implementation	Peláez-Rodríguez et al. [6,39]

Parallel to these developments, Convolutional Neural Networks (CNNs) emerged as a powerful category for fog detection and short-term forecasting, particularly when using input sources such as satellite imagery, sky-view camera feeds, or gridded numerical weather products. CNNs excel at learning spatial features such as fog microstructures, aerosol plumes, and low-level stratus cloud formations. Their ability to extract meaningful visual indicators from large-scale spatial datasets significantly advanced fog classification and near-real-time detection tasks. However, CNN models alone are not inherently temporal, and their forecasting skill deteriorates when long-term atmospheric memory is needed [36,37].

Building upon these complementary strengths, hybrid architectures such as ConvLSTM (Convolutional LSTM) were developed, combining CNN-based spatial feature extraction with LSTM-based temporal learning. These models represent a major advancement in spatiotemporal fog forecasting and have produced state-of-the-art results in applications involving sequential satellite images, NWP visibility fields, or video-based fog monitoring systems [27,28]. Despite achieving superior accuracy, ConvLSTM networks are computationally intensive and require carefully tuned model configurations to ensure stable training.

Another emerging class of models includes Graph Neural Networks (GNNs) and Spatiotemporal Graph Neural Networks (ST-GNNs), which model geographical weather stations and atmospheric fields as interconnected graph structures. These models are especially useful in fog systems such as the Indo-Gangetic Plains, where fog formation and propagation exhibit spatial dependence driven by synoptic weather patterns, cold pool dynamics, and aerosol concentration gradients. Although this approach offers promising spatial generalization capability, challenges remain regarding graph construction rules, scalability, and interpretability [38].

More recent advancements have focused on transformer architectures and attention-based forecasting models, which eliminate the sequential bottleneck in traditional RNNs and enable long-range temporal dependency learning through self-attention mechanisms. Transformers have demonstrated promising results in multi-step visibility prediction and multimodal fusion tasks integrating satellite imagery, synoptic station data, and numerical weather prediction ensembles. While transformers represent a forward-looking direction for operational fog forecasting, they require significantly large training datasets and remain computationally demanding [31–33]. A synthesis of the comparative characteristics of these model families is presented in **Table 2**.

2.3. Bayesian and Probabilistic Modelling in Fog Forecasting

Various researchers have applied different methodologies for prediction of fog [40–44]. While deep learning architectures have demonstrated strong capability in capturing nonlinear atmospheric behaviour, they inherently produce deterministic point predictions. This deterministic nature often limits operational applicability in high-risk domains such as aviation safety, autonomous navigation, and transportation management, where uncertainty surrounding a prediction is equally if not more important than the predicted value itself. Fog forecasting is especially sensitive to uncertainty due to the chaotic nature of boundary-layer thermodynamics, rapid fog onset transitions, and stochastic atmospheric interactions such as aerosol–radiation feedbacks, turbulence fluctuations, and moisture convergence variability [45–48]. These complexities have driven growing interest in Bayesian and probabilistic forecasting frameworks, which explicitly quantify prediction uncertainty rather than producing a single numerical estimate [49–51].

Bayesian approaches estimate the probability distribution of the forecast variable by treating model parameters as random variables and updating their likelihood using observational evidence. Unlike conventional statistical regression where the goal is to find a single optimal model, Bayesian inference produces a posterior distribution that reflects both epistemic uncertainty (due to incomplete knowledge or limited training data) and aleatoric uncertainty (inherent noise in meteorological systems) [52–54]. This probabilistic representation is particularly valuable for fog forecasting because observational datasets are noisy, measurement precision varies across instruments, and atmospheric boundary-layer conditions often exhibit abrupt regime shifts [49,55].

Historically, Bayesian models have been employed in meteorology through probabilistic ensemble prediction systems, Bayesian linear regression, and hierarchical probabilistic modelling frameworks [49,55,56]. However, only limited studies have explored their application to fog prediction due to computational constraints and the difficulty of modelling multimodal atmospheric uncertainty. The emergence of advanced computational tools and variational inference techniques has revived interest in Bayesian modelling as a complement to machine learning and deep learning systems [57–61].

One of the most promising approaches involves Bayesian Neural Networks (BNNs), which incorporate uncertainty estimation into deep learning architectures by replacing fixed weights with probability distributions. During inference, the model samples from the learned distributions, generating a posterior predictive distribution rather than a single deterministic output. Techniques such as Monte Carlo Dropout, Markov Chain Monte Carlo (MCMC), and Variational Bayesian Inference are commonly used to approximate posterior distributions without requiring prohibitively expensive computations [52,58,59]. These methods have demonstrated strong potential in weather and visibility forecasting, particularly at longer lead times where uncertainty grows significantly [50,57].

In addition to Bayesian neural networks, probabilistic forecasting frameworks such as Gaussian Processes (GPs) and Bayesian regression models have been explored in environmental modelling. Gaussian Processes provide smooth uncertainty-aware predictions and are particularly effective when datasets are limited or sparse conditions frequently observed in fog research [55,60]. However, GPs scale poorly with large training datasets and therefore struggle when applied to long-term visibility or fog index forecasting involving decades of continuous meteorological observations [60].

Recent advancements have also focused on Bayesian calibration techniques, which adjust deterministic model outputs to improve confidence interval reliability and reduce overconfidence bias, a known issue in deep learning–based forecasting systems. Calibration strategies such as ensemble learning, posterior confidence scaling, and Bayesian beta regression have demonstrated improved reliability for classification thresholds used in aviation fog warnings, transportation advisories, and early warning dissemination systems [49,57,61].

More contemporary research has emphasized the integration of Bayesian modelling with deep learning architectures to create hybrid systems capable of both accurate forecasting and uncertainty estimation. For instance, studies utilizing Bayesian LSTM models, Bayesian ConvLSTM frameworks, and Bayesian encoder–decoder architectures have shown improvement in confidence-aware forecasting skill [50,51,57]. In particular, hybrid probability models have demonstrated better interpretability when forecasting fog conditions in complex terrains such as coastal fog belts, marine shipping lanes, and inland radiation fog zones including the Indo-Gangetic Plains [50,51].

The relevance of Bayesian modelling for fog forecasting continues to grow due to strategic operational needs. Aviation authorities, rail network operations, and intelligent transport systems increasingly rely on uncertainty-aware predictions to support dynamic decision-making, delay-risk estimation, and hazard-level classification [50,55,56]. The combination of deep learning accuracy and Bayesian interpretability offers a promising direction for next-generation fog forecasting frameworks, especially when extended to forecast variables such as the Fog Index rather than raw visibility, enabling continuous probabilistic environmental risk assessment [50,51,57,61].

2.4. Research Gaps

Despite significant progress in fog detection and short-term visibility prediction using deep learning and probabilistic modelling, several key challenges remain unresolved, particularly for long-term forecasting and operational decision support. The existing body of literature demonstrates substantial advancements in modelling spatial patterns using CNN-based architectures, temporal dependency learning using LSTM and GRU networks, and hybrid multimodal frameworks incorporating satellite observations, reanalysis data, and ground-based meteorological measurements. However, a careful review of current research reveals multiple gaps that limit the reliability, generalizability, and operational readiness of fog forecasting systems, especially for climate-sensitive and high-risk regions such as the IGP.

One of the primary gaps lies in the forecasting horizon. Most studies focus on short-term fog prediction with lead times ranging from 30 min to 12 h. While these short-term models are useful for real-time hazard alerts, they provide insufficient support for strategic planning in aviation scheduling, rail traffic management, agricultural decision-making, and urban transportation operations. Long-term fog prediction (48–120 h), which is operationally more valuable, remains a largely unexplored domain due to increasing uncertainty and model instability over extended time sequences.

A second critical research gap is related to the predictive variable itself. The majority of prior research models raw visibility as the forecasting target. Visibility, however, fluctuates rapidly and does not always reflect the continuity or persistence of fog events. The lack of a standardized, temporally smooth indicator complicates learning and introduces noise-driven instability into forecasting models. The use of a derived metric such as the Fog Index, which incorporates fog intensity and duration into a normalized continuous scale, has received limited attention in existing literature, despite its potential to improve long-term forecast reliability.

A third gap concerns the treatment of uncertainty. Most deep learning models generate deterministic (single-valued) predictions, offering no information regarding prediction confidence or uncertainty distribution. In domains such as aviation or intelligent transportation systems, uncertainty quantification is essential for risk-sensitive decision-making and operational constraints management. Although Bayesian and probabilistic approaches exist, they have not been systematically integrated with deep sequential learning architectures for fog forecasting. Current probabilistic models also lack calibration mechanisms tuned for meteorological systems with non-linear seasonal structures and high spatiotemporal variability.

Another limitation identified in existing research is the lack of comprehensive spatiotemporal hybrid modelling frameworks capable of integrating multivariate atmospheric dynamics with probabilistic inference. While ConvLSTM, transformers, and hybrid CNN-RNN models have shown success for spatiotemporal learning, they rarely address uncertainty quantification, interpretability, or long-horizon forecasting capability. Similarly, Bayesian neural networks and Gaussian processes address probabilistic forecasting but struggle with high-dimensional temporal dependencies and complex atmospheric evolution patterns.

A further gap relates to dataset diversity and operational validation. Although datasets from European coastal fog belts, marine fog systems, and satellite fog mask products are widely studied, there is limited research focused specifically on the IGP a region with some of the highest fog frequency globally and unique boundary-layer behaviour driven by aerosol loading, agricultural emissions, and Himalayan cold air drainage. Additionally, few studies validate fog forecast performance across multiple seasonal cycles or multi-decade datasets, which is necessary to ensure model robustness in climate-variable environments.

Taken together, these observations highlight the need for a hybrid forecasting framework capable of (a) learning long-term fog dynamics using deep temporal models, (b) generating interpretable probabilistic outputs, and (c) forecasting a meaningful, stable prediction variable such as the Fog Index rather than raw visibility.

2.5. Problem Statement

Fog forecasting over the IGP remains operationally challenging due to strong nonlinearity, data variability, and uncertainty in atmospheric processes governing fog onset, persistence, and dissipation. Existing forecasting systems predominantly offer short-term deterministic predictions, rely on raw visibility measurements, and lack mechanisms to quantify forecast uncertainty limiting their usability for high-risk operational environments such as aviation, ground transportation, and early warning systems.

Therefore, there is a compelling need to develop an integrated forecasting model that:

- Can predict fog conditions reliably over extended lead times (up to five days).
- Uses a stable, derived environmental metric such as the Fog Index instead of raw visibility.
- Combines deep temporal learning capability with Bayesian probabilistic inference.
- Provides prediction intervals and uncertainty metrics suitable for operational risk assessment.

3. Methodology

The methodology adopted for developing the proposed hybrid long-term fog forecasting framework includes a three-layer BiLSTM architecture integrated with a Modified Bayesian Beta Regression layer. The methodology is structured to ensure scientific rigor, operational applicability, and reproducibility. It includes dataset acquisition, preprocessing, feature engineering, Fog Index computation, deep learning architecture design, Bayesian uncertainty modelling, training–validation strategy, and evaluation metrics. The overall framework follows a systematic pipeline in which raw meteorological observations are converted into structured time-series sequences, normalized and transformed to enhance predictive signals, and subsequently processed through the hybrid neural-Bayesian forecasting model. The methodology also incorporates multiple enhancements, including Monte-Carlo dropout for probabilistic inference, hyperparameter optimization, and model calibration to ensure robust performance over extended forecast horizons.

3.1. Problem Definition and Modelling Framework

Deep learning architectures such as BiLSTM networks and probabilistic frameworks such as Bayesian Beta Regression have each demonstrated strong potential in fog and visibility forecasting. However, each approach has inherent limitations when applied independently. BiLSTM networks are highly effective in learning temporal evolution patterns from sequential atmospheric data and can capture nonlinear dependencies, multi-step lag effects, and long-term seasonal behaviour. Despite this strength, they produce deterministic point forecasts without expressing how confident the model is in its predictions an aspect critically important for operational sectors such as aviation safety, highway traffic control, and railway scheduling. In contrast, Bayesian Beta Regression provides a statistically grounded mechanism for forecasting with uncertainty representation through probabilistic distributions and credible intervals, yet it lacks the ability to autonomously learn complex time-dependent patterns embedded in meteorological sequences.

This complementary weakness creates a methodological gap between models that are accurate but not uncertainty-aware, and those that are uncertainty-aware but not temporally expressive. The objective of this study is to address this gap by constructing a hybrid forecasting framework that integrates a three-layer BiLSTM network with a Modified Bayesian Beta Regression layer. In the proposed system, the BiLSTM component first learns the temporal structure of fog behaviour from historical meteorological records, while the Bayesian regression layer transforms the learned latent temporal representation into a probabilistic forecast of the Fog Index. The result is a forecasting system capable of generating precise numerical predictions while simultaneously quantifying predictive uncertainty, thereby enabling risk-aware and operationally meaningful decision support in fog-prone environments. Formally, fog prediction in this study is formulated as a multi-step sequence-to-sequence regression task. Let:

$$X_t = \left(X_t^1, X_t^2, \ldots, X_t^n\right) \tag{13}$$

represent the vector of meteorological variables at time step

t, where values may include atmospheric temperature, dew-point temperature, relative humidity, pressure, visibility, wind speed, and derived metrics such as Fog Index. Given an input window length L, the historical input sequence is written as:

$$X_{(t)} = \left[X_{(t-L+1)}, X_{(t-L+2),} \ldots X_{(t)}\right] \quad (14)$$

The forecasting objective is to learn a nonlinear function $f(\theta)$ that predicts the Fog Index at future time horizons. For a single future lead-time Δ, the prediction is represented as:

$$\widehat{FI}(t+\Delta) = f_{(\theta)} X_{(t)} \quad (15)$$

In this study, the model predicts multiple future Fog Index values corresponding to several lead times including 6 h, 1 day, 3 days, and 5 days. Therefore, the prediction generalizes to:

$$\begin{gathered}\left[\widehat{FI}\left(t+\Delta_1\right), \widehat{FI}\left(t+\Delta_2\right), \ldots, \widehat{FI}\left(t+\Delta_m\right)\right] \\ = f(\theta)\left(X(t)\right)\end{gathered} \quad (16)$$

Within the hybrid architecture, the model function $f(\theta)$ is decomposed into two components. First, a three-layer Bidirectional LSTM encoder, represented as:

$$h(t) = g(\varphi)\left(X(t)\right) \quad (17)$$

learns temporal dependencies and produces a latent representation of the input sequence. Second, a Modified Bayesian Beta Regression layer interprets this representation and models the Fog Index at each forecast step as a Beta-distributed random variable:

$$FI(t+\Delta) \mid h(t) \sim \text{Beta}\left(\alpha\left(h(t)\right), \beta\left(h(t)\right)\right) \quad (18)$$

where $\alpha(h(t))$ and $\beta(h(t))$ are positive-valued outputs predicted by the Bayesian regression layer.

The predictive mean of the Fog Index is computed as:

$$\text{Mean} = \frac{\alpha}{\alpha+\beta} \quad (19)$$

and the predictive variance is computed as:

$$\text{Variance} = \frac{\alpha\beta}{(\alpha+\beta)^2\left(\alpha+\beta+1\right)} \quad (20)$$

The variance defines the uncertainty width of the prediction interval; a higher variance indicates lower confidence, while a lower variance indicates a more certain prediction.

3.2. Methodological Overview

The methodological workflow adopted in this study is designed to ensure systematic transformation of raw meteorological observations into reliable long-term fog forecasts enriched with quantified uncertainty. The process integrates classical data engineering techniques with advanced deep learning and probabilistic modelling approaches. The workflow consists of seven major stages, as described in detail below.

3.2.1. Data Acquisition

The forecasting framework developed in this study is based on a comprehensive 24-year meteorological dataset spanning January 2000 to December 2023, collected from Lucknow Airport an important location within the IGP, where dense winter fog events frequently occur. The dataset consists of multiple atmospheric variables recorded at 6-h intervals, allowing the model to capture essential transitions in fog evolution, including onset, peak intensity near sunrise, and dissipation phases, while minimizing noise from excessively fine temporal sampling. The records form a structured multivariate sequence mathematically defined as:

$$X = \{\, x_t \mid x_t \in R^n,\, t = 1,\, 2,\, \ldots,\, T \,\} \quad (21)$$

Each observation can be expressed as a feature vector of atmospheric variables relevant to fog formation:

$$\begin{gathered} x_t = \\ \left[\, T_t,\, Td_t,\, RH_t,\, Ws_t,\, Ps_t,\, Vis_t,\, f_{derive}(t)\right]\end{gathered} \quad (22)$$

where T_t is air temperature, Td_t is dew-point temperature, RH_t is relative humidity, Ws_t is wind speed, Ps_t is surface pressure, and Vis_t represents measured visibility. The term $f_{derived}(t)$ includes additional computed predictors such as dew-point depression ($T_t - Td_t$), vapour pressure, or air-mass moisture indicators, which enhance feature representation for deep learning. The long temporal duration and rich multivariate structure of this dataset enable the model to learn both seasonal fog dynamics and inter-annual atmospheric variability influenced by aerosols, radiative cooling, synoptic forcing, and land-atmosphere interactions. This makes the dataset highly suitable for developing a hybrid deep learning and Bayesian forecasting architecture capable of supporting operational fog prediction and uncertainty estimation.

3.2.2. Data Cleaning and Preprocessing

Data preprocessing is a necessary step to ensure that the forecasting model receives consistent, accurate, and well-structured input. Meteorological time-series data typically include missing values, sudden erroneous spikes, temporal discontinuities, and reporting inconsistencies caused by instrument malfunction, atmospheric noise, or logging delays. Therefore, before model training, the raw dataset underwent several processing steps including missing value correction, outlier detection, temporal smoothing, normalization, and dataset restructuring into supervised learning format.

Missing observations were first detected by identifying timestamps with incomplete or null entries. Short gaps were handled using linear interpolation, which estimates intermediate values along a straight line between two known points. The interpolation formula used was:

$$x(t) = x\left(t_1\right) + \left(x\left(t_2\right) - x\left(t_1\right)\right)/(t_2 - t_1) \cdot (t - t_1) \quad (23)$$

where t_1 and t_2 denote neighbouring timestamps surrounding the missing value.

For larger gaps (more than one interval), Exponential Weighted Moving Average (EWMA) smoothing was applied to preserve temporal fog dynamics without introducing unrealistically sharp transitions. The smoothing was performed using:

$$x_s(t) = \alpha\, x(t) + (1 - \alpha)\, x_s(t-1) \quad (24)$$

where α (alpha) is the smoothing factor, set between 0.4 and 0.6 depending on variable sensitivity.

Outlier detection was performed using z-score filtering, which identifies statistically abnormal values by comparing each observation to the dataset mean and standard deviation. A value was flagged as an outlier if:

$$z = \frac{x - \mu}{\sigma} > 3 \quad (25)$$

where μ is the mean and σ is the standard deviation of the respective feature. Outliers were replaced either by the nearest non-anomalous entry or by the median of the rolling window, depending on whether the anomaly was a brief spike or part of a measurement fault.

After correcting missing and anomalous values, temporal smoothing was applied using a rolling average method to reduce sensor noise while preserving diurnal fog evolution patterns. The rolling average was computed as:

$$\overline{x}(t) = \frac{1}{k} \sum_{i=0}^{k-1} x(t - i) \quad (26)$$

where k = 3 was chosen based on empirical evaluation to prevent excessive smoothing.

Since deep learning models are highly sensitive to feature scale differences, all variables were normalized using Min-Max scaling, which transforms each feature into a homogeneous numeric range between 0 and 1, defined as:

$$x^{'} = \frac{x - x_{\min}}{x_{\max} - x_{\min}} \quad (27)$$

This scaling ensures faster convergence during model training and prevents dominant influence from high-magnitude variables such as pressure or visibility.

Finally, to prepare the data for time-series forecasting, the continuous meteorological records were converted into supervised learning format, where sequential input windows were paired with future prediction targets. If L represents the sequence length (window size), then model input and output pairs were generated as:

$$\text{Input:} \sim [\, x(t-L),\, x(t-L+1),\, \ldots,\, x(t)\,]$$

$$\text{Output:} \sim \text{Fog Index}\,(t + \text{forecast lead time})$$

This transformation allowed the deep learning model to learn temporal dependencies required for multi-step forecasting.

3.2.3. Fog Index Construction

Visibility is one of the most widely used meteorological parameters for quantifying fog severity; however, the raw visibility values tend to fluctuate sharply due to atmospheric turbulence, sensor measurement inconsistencies, and rapidly changing near-surface radiative conditions. These fluctuations introduce significant noise into forecasting models and reduce predictive stability, especially when applied to multi-step forecasting tasks. To address this issue, the present study transforms visibility observations into a normalized and dimensionless Fog Index (FI) that represents fog severity on a continuous scale from 0 to 1. In this scale, a value closer to zero corresponds to clear atmospheric conditions, while values approaching one represent dense or very dense fog. This formulation produces a smoother temporal signal, better supports long-term sequence learning architectures such as BiLSTM, and facilitates uncertainty quantification in the Bayesian regression layer.

The Fog Index formulation is conceptually linked to aviation category thresholds used by the India Meteorological Department (IMD) and International Civil Aviation Organization (ICAO), where visibility greater than 1,000 m corresponds to no operational fog restrictions, moderate fog occurs between 200 and 500 m, and dense fog conditions are typically observed when visibility falls below 200 m. While such categorical boundaries are useful in operational meteorology, deep learning models benefit from a continuous representation that captures gradual transitions rather than abrupt category shifts. Therefore, visibility measurements are mathematically transformed using a normalized inverse-scaling approach, defined as:

$$FI(t) = 1 - \frac{Vis(t) - Vis_{\min}}{Vis_{\max} - Vis_{\min}} \tag{28}$$

where $FI(t)$ denotes the Fog Index at time t, $Vis(t)$ is the observed visibility, and $Vis_{\min}$ and $Vis_{\max}$ represent the minimum and maximum visibility values within the dataset. This transformation ensures the expected monotonic behaviour: as visibility decreases, the Fog Index increases.

Since raw meteorological data, including visibility, often exhibit short-term spikes inconsistent with typical fog lifecycle patterns, smoothing is applied to enhance temporal coherence. A rolling mean function is used to minimize transient variations while preserving the general trend of fog progression. The smoothed signal is computed as:

$$FI_s(t) = \frac{1}{k} \sum_{i=0}^{k-1} FI(t-i) \tag{29}$$

where $k = 3$ defines the window length based on empirical testing. This smoothing process produces a more stable and interpretable Fog Index curve over time.

Fog behaviour is additionally characterized by persistence; once formed, fog rarely dissipates instantaneously. To model this physical behaviour, a persistence adjustment term is introduced using an exponential update mechanism:

$$FI^*(t) = \lambda\, FI(t) + (1-\lambda)\, FI^*(t-1) \tag{30}$$

where λ is a weighting factor tuned between 0.5 and 0.7 to balance sensitivity and inertia. This recursive formulation ensures that the Fog Index evolves smoothly over successive timestamps, improving the temporal realism of the modelled fog lifecycle.

In its final form, the Fog Index used in this study represents a processed, normalized, and physically consistent signal derived from raw visibility observations. It retains essential meteorological information regarding fog intensity and persistence while reducing noise and improving suitability for multi-step deep learning forecasting. The Fog Index therefore becomes not only a prediction target, but also an interpretive tool that bridges raw observational data and probabilistic model output, supporting operational forecasting applications in aviation, road transport, and public safety.

To facilitate reproducibility, the Fog Index calculation follows a three-stage pipeline: (a) Inverse scaling of raw visibility (Equation (28)), (b) Rolling mean smoothing to remove noise (Equation (29)), and (c) Exponential update for physical persistence (Equation (30)). This stabilizes the target signal for long-term forecasting compared to raw visibility in meters. Further, the Fog Index improves forecasting over raw visibility because: (a) It reduces high-frequency noise, (b) Captures persistence of fog events, and (c) Provides bounded learning space (0–1). This leads to more stable long-term predictions compared to highly volatile visibility values.

3.2.4. Feature Engineering and Transformation

Feature engineering plays a critical role in improving model learning efficiency by enhancing signal quality and exposing hidden structures within atmospheric data. The meteorological variables included in this study possess different physical units, ranges, and variabilities. For example, air temperature may vary between 0 °C and 40 °C, while relative humidity ranges from 0–100%, wind speed may fluctuate between 0–10 m/s, and visibility can change rapidly from a few meters during dense fog to several kilometers under clear conditions. Without scaling, such disparities may cause the deep learning model to overweight high-magnitude variables such as pressure or visibility while underrepresenting equally important but numerically smaller features. To resolve this imbalance and ensure consistent gradient behaviour during training, feature values were normalized using Min-Max scaling. The transformation is mathematically defined as:

$$x' = \frac{x - x_{\min}}{x_{\max} - x_{\min}} \tag{31}$$

In this transformation, x' represents the normalized feature value, x denotes the original observed value, and $x_{\min}$ and $x_{\max}$ correspond to the minimum and maximum values of that feature across the entire dataset. By applying this scaling, all atmospheric variables are compressed into a uniform range between 0 and 1. This normalization step ensures balanced feature contribution during training and prevents variables with naturally larger numerical magnitudes, such as visibility or pressure from dominating the learning process. As a result, the model is able to interpret each feature proportionately, improving optimization stability, accelerating convergence, and enhancing the overall robustness of the forecasting framework.

Beyond normalization, capturing temporal structure within the fog dataset is essential because fog occurrence in the IGP exhibits strong seasonal and periodic recurrence, especially during December and January when synoptic conditions, radiative cooling, and moisture availability align. These seasonal signals may not be fully captured by raw time-series patterns alone; therefore, spectral decomposition techniques were applied to extract cyclical components embedded within the data. To achieve this, the Fast Fourier Transform (FFT) was applied to selected meteorological features such as humidity, dew-point depression, and visibility, converting the time-domain signals into corresponding frequency-domain representations. FFT identifies the dominant periodicities and seasonal harmonics that contribute to fog behaviour. The transformation is formally expressed as:

$$F(k) = \sum_{t=0}^{T-1} x(t)\, e^{-i2\pi kt/T} \tag{32}$$

In this equation, $F(k)$ represents the complex spectral coefficient at frequency k, while $x(t)$ denotes the original time-series signal being transformed. The term T corresponds to the total number of observations in the dataset, and the exponential component $e^{-i2\pi kt/T}$ serves as the complex sinusoidal wave basis used in Fourier decomposition, enabling the conversion of temporal meteorological variations into their underlying frequency components.

The magnitude and phase of the resulting spectral components were analysed to identify dominant low-frequency patterns associated with winter fog cycles. The most relevant harmonics were then incorporated into the model as auxiliary features, allowing the deep learning architecture to learn not only short-term fog transitions, but also broader recurring climatological patterns influenced by annual cold season dynamics, aerosol accumulation, and boundary-layer stability.

By combining normalization and spectral feature expansion, the input space becomes more informative and structured, enabling the BiLSTM-Bayesian hybrid model to better detect relationships among meteorological drivers, seasonal cycles, and fog intensity behaviour. This preprocessing ensures improved model generalization, faster convergence, and enhanced forecasting accuracy, particularly for multi-step fog predictions requiring recognition of long-term seasonal dependencies.

Although the proposed framework is designed to incorporate multivariate meteorological inputs, including temperature, humidity, dew-point, and FFT-derived spectral features, the present study evaluates the model using the Fog Index as a univariate input. This design choice was made to isolate the intrinsic temporal learning capability of the BiLSTM architecture and to maintain consistency with prior benchmark studies employing the same Fog Index formulation. Consequently, the multivariate preprocessing pipeline is presented as part of the overall framework capability and is intended for future extension of the model.

3.2.5. Deep Learning Sequence Modelling Using BiLSTM

Fog formation is inherently a temporal phenomenon influenced by gradual atmospheric transitions rather than isolated observations. Factors such as radiative cooling, high moisture accumulation, declining dew point depression, and stagnant wind flow typically evolve over several hours before dense fog materializes. Similarly, fog dissipation follows progressive warming and increased atmospheric mixing rather than abrupt atmospheric shifts. Therefore, modelling fog dynamics requires a framework capable of capturing long-term sequential dependencies, rather than relying on independent or shallow statistical relationships. Deep learning architectures based on RNNs are well-suited for this task because they are designed to process time-dependent inputs; however, standard RNNs suffer from vanishing gradient limitations and struggle to retain information across long temporal windows. To overcome these issues, this study employs a BiLSTM architecture, which improves representational capability by learning temporal patterns in both forward and backward chronological directions.

The LSTM mechanism extends standard RNNs by incorporating gating functions that regulate how information is stored, forgotten, and propagated over time. This enables the model to retain memory of relevant meteorological states over many time steps, making it highly suitable for fog forecasting where atmospheric conditions evolve slowly but meaningfully. A three-layer BiLSTM network was selected after comparative testing to provide sufficient depth for hierarchical temporal feature extraction without overfitting. Unlike a unidirectional LSTM, which processes the input sequence from past to future, a BiLSTM performs dual processing: one pass moves forward through time while the other moves backward. This dual processing provides additional context because fog evolution may depend not only on preceding atmospheric conditions but also on patterns that become apparent only when the full temporal behaviour is considered. Mathematically, for each time step *t*, the forward LSTM cell computes its hidden state based on the current input and the previous forward state:

$$\overrightarrow{h}(t) = \text{LSTM}\left(x(t),\ \overrightarrow{h}(t-1)\right) \tag{33}$$

Simultaneously, a backward LSTM processes the sequence in reverse order and generates a hidden state based on the upcoming observation rather than the previous one:

$$h(t)\ \leftarrow\ \text{LSTM}\left(x(t),\ \text{h}(t+1)\right) \tag{34}$$

Once both directional states are computed, they are merged to form a richer contextual representation. The combined hidden representation at time step *t* is formed by concatenating the forward and backward states:

$$h(t) = \left[\, h_{\rightarrow}(t)\ \oplus\ h_{\leftarrow}(t) \,\right] \tag{35}$$

This concatenated representation ensures that the model can learn relationships where future values improve understanding of present patterns an important advantage for fog prediction, where dissipation patterns and repetition cycles inform onset prediction and uncertainty estimation.

The selection of the three-layer BiLSTM architecture was guided not only by empirical performance but also by the multi-scale temporal nature of fog evolution. Fog formation and persistence are influenced by interacting processes operating across different temporal scales, including short-term micro-meteorological adjustments, diurnal radiative cycles, and multi-day boundary-layer stability regimes. Shallow recurrent architectures (one or two layers) were insufficient to simultaneously capture both rapid fluctuations and longer persistence behaviour. In contrast, deeper architectures exceeding three layers increased model variance and susceptibility to overfitting without providing meaningful improvements in validation performance. BiLSTM was selected because: (a) It works well with moderate dataset size; (b) Lowers computational cost than transformers; and (c) Superior performance observed experimentally. Transformers showed overfitting due to limited data.

The three-layer configuration therefore offers a practical balance between temporal representational capacity and generalization stability. Within this structure, the first layer primarily captures short-term meteorological variability, while subsequent layers progressively abstract persistence-related patterns associated with fog formation and maintenance. This hierarchical temporal learning is particularly important for multi-day forecasting tasks, where fog evolution depends on cumulative atmospheric memory rather than instantaneous conditions.

The bidirectional design further enhances the model's ability to recognize patterns related to fog onset, peak intensity, and dissipation by capturing both historical dependencies (e.g., gradual humidity buildup) and forward contextual signals (e.g., persistence of inversion layers or wind stagnation). Compared to unidirectional LSTM architectures, the BiLSTM provides improved representation of meteorological processes governed by boundary-layer dynamics. In this study, the BiLSTM functions as the core temporal learning engine, transforming raw meteorological history into latent representations that are subsequently passed to the Bayesian regression layer for probabilistic forecasting. This hybrid integration enables not only accurate estimation of future fog severity but also quantification of predictive uncertainty, making the framework suitable for operational early-warning systems, aviation scheduling, transportation planning, and other risk-sensitive applications in fog-prone environments.

3.2.6. Modified Bayesian Regression for Uncertainty Quantification

While the BiLSTM component provides accurate sequence learning capability, its output is inherently deterministic and does not convey the uncertainty associated with predictions. Fog forecasting, however, involves inherently stochastic atmospheric behaviour influenced by nonlinear interactions such as aerosol–radiation feedbacks, boundary-layer turbulence, synoptic pressure variation, and anthropogenic emissions. Because of these uncertain and often unpredictable influences, two forecasts with identical input conditions may produce different fog outcomes. Therefore, meaningful operational fog prediction requires not only a point estimate of the Fog Index but also a quantifiable measure of confidence surrounding that estimate. To address this requirement, the deterministic output produced by the final BiLSTM layer is passed into a Modified Bayesian Beta Regression module, enabling probabilistic forecasting and uncertainty representation.

The Fog Index is a bounded variable confined to a continuous range between 0 and 1, making it unsuitable for ordinary Gaussian-based regression models. Instead, the Beta distribution is mathematically appropriate because it supports flexible shapes, asymmetry, and bounded intervals-properties that align with the behaviour of fog intensity. In principle, alternative likelihood formulations such as a Gaussian model combined with a logistic transformation could also be used to constrain predictions within the (0,1) interval. However, such approaches assume symmetric error behaviour in the transformed space and may distort uncertainty near the boundaries. Truncated Gaussian models similarly lack flexibility in representing skewness and saturation effects commonly observed during fog formation and persistence phases. In contrast, the Beta distribution naturally supports bounded continuous variables and allows asymmetric shapes without requiring transformation. Given that the Fog Index exhibits regime-dependent skewness and saturation behaviour, the Beta likelihood provides a more physically consistent probabilistic representation. Thus, the forecast target is modelled as:

$$FI(t) \sim \mathrm{Beta}(\alpha,\ \beta) \tag{36}$$

where α and β are shape parameters learned during the regression process. The expected value (mean) and variance of the Beta distribution, which define the central tendency and uncertainty of the probabilistic forecast.

To approximate Bayesian posterior inference without requiring computationally expensive sampling methods such as Markov Chain Monte Carlo (MCMC), the model incorporates Monte-Carlo Dropout, a practical approximation that introduces stochasticity during inference. Instead of disabling dropout after training as done in standard neural networks, the dropout mechanism remains active across repeated forward passes. Each pass perturbs the model weights, enabling sampling from an implicit posterior distribution rather than generating a single deterministic output. The final probabilistic forecast is computed as the mean of repeated forward passes:

$$\widehat{y} = \frac{1}{N}\sum_{i=1}^{N} f\left(x,\ \theta_i\right) \tag{37}$$

where θ_i represents a randomly sampled dropout-perturbed parameter state and N = 100 forward passes were used in this study to generate a stable posterior estimate. This sampling strategy produces not only a mean prediction but also a distribution of possible outcomes from which credible intervals, such as 50%, 80%, or 95% confidence bands—can be computed.

By combining deep temporal learning (BiLSTM) with a stochastic probabilistic inference mechanism (Bayesian Beta Regression with Monte-Carlo Dropout), the proposed hybrid model is able to generate fog forecasts that reflect both accuracy and reliability. This uncertainty-aware design is particularly important in aviation scheduling, railway dispatching, and intelligent transportation systems, where operational decisions must weigh the likelihood of severe fog conditions rather than rely solely on deterministic predictions.

The following three modifications were introduced: (a) Neural Parameterization: α and β are predicted from BiLSTM outputs instead of linear predictors; (b) Monte Carlo Dropout Integration: Enables approximate Bayesian inference; and (c) Temporal Dependency Embedding: Regression operates on sequence-learned latent features instead of static inputs. Aleatoric uncertainty is captured via Beta distribution variance. Epistemic uncertainty is estimated using Monte Carlo dropout.

Although classical probabilistic models such as Quantile Regression Forests and Gaussian likelihood-based approaches could be considered as alternative baselines, their applicability is limited in the present context. These models do not explicitly capture temporal dependencies inherent in fog evolution and often assume symmetric error behaviour. Furthermore, Gaussian-based formulations require transformation to enforce bounded outputs, which may distort uncertainty near the boundaries. In contrast, the Beta distribution naturally supports bounded continuous variables in the (0,1) range and provides greater flexibility in modelling skewness and saturation effects observed in fog dynamics.

3.2.7. Model Evaluation and Calibration

To ensure objective and reliable performance assessment, the dataset was partitioned using a temporal splitting strategy rather than random sampling. This preserves chronological order and prevents information leakage from future observations into past training sequences. The dataset (2000–2023) was divided into three non-overlapping periods: training (2000–2016), validation (2017–2019), and testing (2020–2023). The test set was strictly placed after the training and validation periods so that model evaluation was conducted only on future unseen observations, thereby maintaining a realistic forecasting setup. Model accuracy was quantified using standard statistical measures widely applied in regression-based environmental forecasting studies. The Root Mean Squared Error (RMSE) was used to measure the magnitude of prediction deviations:

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \widehat{y}_i)^2} \tag{38}$$

where y_i denotes the observed Fog Index and $\hat{y}_i$ denotes the predicted value at time step i. To complement this, the Mean Absolute Error (MAE) was calculated to capture the average error magnitude without disproportionately penalizing outliers:

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \widehat{y}_i| \tag{39}$$

Since the proposed framework produces probabilistic predictions through Bayesian inference, additional metrics were used to evaluate uncertainty quality. The Continuous Ranked Probability Score (CRPS) was employed to assess probabilistic sharpness and calibration. Furthermore, the Prediction Coverage Probability (PCP) was computed to assess the frequency with which the observed Fog Index values were contained within the anticipated confidence intervals. A well-calibrated uncertainty model should yield PCP values close to the nominal confidence level (e.g., ~80% for an 80% interval).

3.3. System Architecture

The development of the proposed hybrid long-term fog forecasting model required the integration of multiple computational components arranged in a structured pipeline designed to support data processing, sequential learning, probabilistic inference, and prediction evaluation. The resulting architecture represents a fusion of classical statistical processing, deep learning time-series modelling, and Bayesian uncertainty quantification, allowing the system not only to forecast fog severity but also to provide estimates of model confidence. The architecture is organized into a series of interdependent layers, each performing a specialized function that contributes to the overall forecasting capability.

The first component of the system is the data processing layer, responsible for preparing raw meteorological observations for modelling. This layer performs essential preprocessing tasks including dataset ingestion, temporal alignment, missing value correction, outlier handling, scaling through Min-Max normalization, and the construction of the Fog Index. By transforming the highly variable raw visibility measurements into a smooth, bounded index between 0 and 1, the system ensures that the subsequent layers receive consistent, normalized, and meaningful predictor values. This step is especially critical in atmospheric applications, where physical measurements span different numerical scales and may contain noise or irregular temporal patterns.

Once processing is complete, the dataset enters the sequential learning layer, where temporal dependencies are modelled using a three-layer BiLSTM network. Unlike conventional feed-forward neural networks, which treat data points independently, BiLSTM networks are explicitly designed to process sequential information by retaining memory of past observations and, in the bidirectional configuration, also incorporating future context. This allows the model to learn the evolving fog lifecycle onset, duration, peak severity, and dissipation patterns based on meteorological behaviour over multiple prior time steps. The output of this layer is a sequence of learned hidden-state representations that capture the temporal evolution of fog-related atmospheric drivers.

The output from the BiLSTM network is then passed to the probabilistic modelling layer, which implements a Modified Bayesian Beta Regression framework. This layer transforms the deterministic output of the deep learning model into a predictive probability distribution rather than a single numerical forecast. The use of a Beta distribution is appropriate because the Fog Index is bounded between 0 and 1. Monte-Carlo Dropout sampling is applied at inference time to approximate Bayesian posterior distributions, enabling the system to quantify uncertainty and generate confidence intervals around each forecast value. This probabilistic perspective is crucial for applications where decision-making depends not only on the expected state of fog but also on the reliability of that estimate, such as in aviation scheduling, highway traffic control, or railway routing during winter fog events.

The final stage of the architecture is the prediction and evaluation layer, which generates the final Fog Index forecast and computes confidence bounds based on the posterior samples. This layer also handles performance scoring using deterministic and probabilistic evaluation metrics, including RMSE, MAE, R^2, the Continuous Ranked Probability Score (CRPS), and PCP. These metrics help quantify both forecasting accuracy and uncertainty calibration, ensuring that the system provides trustworthy outputs suitable for operational decision support in fog-prone environments. A conceptual representation of this end-to-end forecasting framework is shown in **Figure 3**, illustrating how the data flows through the system—from raw meteorological measurements to final uncertainty-aware prediction outputs.

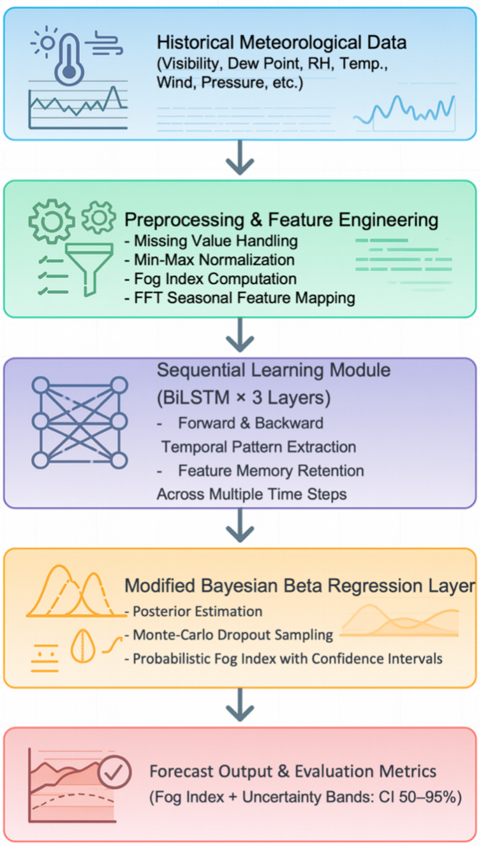

Figure 3. Schematic representation of the proposed hybrid fog forecasting framework integrating a three-layer BiLSTM for temporal sequence learning with a Modified Bayesian Beta Regression layer for probabilistic prediction and uncertainty quantification.

3.4. Rationale for Hybrid Architecture

The choice of a hybrid deep learning and Bayesian modelling framework was driven by both theoretical and operational requirements associated with long-term fog prediction. Fog formation is a complex atmospheric process governed by nonlinear interactions among meteorological variables such as humidity, dew point temperature, wind speed, radiative cooling, boundary-layer inversion strength, and aerosol concentration. These interactions evolve gradually over time, often exhibiting lag effects and seasonal memory patterns. Traditional forecasting models including statistical regression, autoregressive moving average structures, and simple machine learning approaches struggle to accurately capture these temporal dependencies because they assume linear relationships or rely solely on instantaneous patterns without temporal context. Deep learning models, particularly those based on LSTM architectures, have demonstrated strong capability in modelling nonlinearity and learning long-range temporal dependencies. The BiLSTM enhances this capability by processing sequences in both forward and backward directions, allowing the model to understand fog evolution not only based on prior meteorological conditions but also on emerging future trends within the sequence window. This results in a richer and more robust feature representation that reflects real atmospheric dynamics.

However, while deep learning improves prediction accuracy, its limitation lies in the deterministic nature of its outputs. A typical LSTM-based model produces a single numeric forecast for each time step, without offering insight into the confidence or uncertainty associated with the prediction. Fog forecasting is inherently uncertain due to stochastic influences such as localized turbulence, sensor noise, land-use variability, and sudden synoptic-scale weather shifts. In real operational contexts particularly in aviation, railway traffic management, highway safety operations, and automated decision-support systems. It is not sufficient to predict fog severity alone; stakeholders must also understand the reliability of the prediction. For example, a forecast indicating severe fog with low uncertainty requires a different response than a similar prediction accompanied by a wide confidence interval. Pure deep learning methods cannot address this uncertainty dimension, which limits their suitability for decision-critical meteorological forecasting.

Conversely, Bayesian modelling approaches are inherently probabilistic, providing a statistical framework in which model parameters and predictions are represented not as fixed values but as probability distributions. Bayesian regression produces credible intervals, predictive uncertainty bounds, and posterior distributions that help quantify how confident the system is in a generated forecast. Yet, Bayesian methods in isolation lack the structural capacity to learn complex sequential temporal patterns embedded within atmospheric time series. They perform well for uncertainty estimation but cannot independently capture multivariate fog evolution dynamics with the same fidelity as deep neural sequence learning architectures.

Therefore, a hybrid approach that combines both modelling paradigms was adopted to leverage the strengths and resolve the shortcomings of each. In the proposed system, the BiLSTM component functions as the temporal pattern extractor, learning fog evolution signatures, diurnal fluctuations, and slowly varying seasonal influences from the input meteorological sequences. The output of this deep learning stage is then passed into a Modified Bayesian Beta Regression layer, which serves as the probabilistic inference engine. This final layer transforms deterministic outputs into probability distributions, generating prediction intervals and uncertainty-aware outputs that support risk-informed operational decisions.

By integrating BiLSTM-based sequence learning with Bayesian uncertainty modelling, the resulting framework is not only capable of generating accurate fog forecasts but also of representing the confidence associated with those predictions. This hybrid design makes the forecasting system scientifically rigorous, interpretable, and operationally meaningful particularly for applications where forecasting reliability is as critical as accuracy.

3.5. Training Strategy and Hyperparameter Optimization

The training of the proposed deep learning-based fog forecasting models was conducted in a systematic yet practically oriented manner using the Google Colab environment. Instead of relying on automated hyperparameter search methods such as Bayesian optimization or exhaustive grid search, the model configuration was refined iteratively through a series of controlled experiments. These experiments were guided by validation performance, residual behaviour, and the stability of training and validation loss curves. Multiple recurrent architectures were implemented and evaluated under a consistent input-output framework derived from the Fog Index time series. The tested configurations included a vanilla LSTM, stacked LSTM, stacked LSTM with dropout, Bidirectional LSTM, a three-layer Bidirectional LSTM, and a CNN-LSTM hybrid. A comparative assessment of these architectures revealed that single-layer and shallow stacked models were able to capture short-term variability but struggled to represent multi-day persistence patterns characteristic of fog evolution. Increasing the depth beyond three recurrent layers did not yield consistent improvements; instead, it resulted in greater validation loss variability, indicating potential overfitting and reduced generalization capability.

The three-layer BiLSTM configuration was therefore selected as it provided the most stable validation performance across forecasting horizons while maintaining computational efficiency. The use of 64 units in the first recurrent layer was intended to capture high-dimensional short-term atmospheric variability, while subsequent layers with fewer units progressively compressed temporal representations to enhance stability and prevent over-parameterization. Preliminary experiments with residual connections in deeper recurrent stacks were also explored but did not produce measurable performance gains relative to the increased architectural complexity. This suggests that, for the present fog forecasting task, hierarchical temporal representation was more beneficial than increased depth.

For all deep learning models, the dataset was transformed into a supervised learning format using a sliding-window approach. In this setup, 20 past time steps were used to predict the subsequent 20 time steps of the Fog Index at 6-h resolution. Accordingly, the input sequences were structured as

$$X \in R^{N \times 20 \times 1} \tag{40}$$

and the output sequences:

$$y \in R^{N \times 20} \tag{41}$$

where N is the number of samples. This design allowed the model to predict a full 5-day sequence (20 steps × 6 h) from the preceding 5 days of data, and specific lead times such as 6 h, 1 day, 3 days, and 5 days were later extracted from the 20-step output vector.

All models were trained using the Adam optimizer with its default learning rate (as provided by Keras), and the loss function was Mean Squared Error (MSE). In addition, the Root Mean Squared Error (RMSE) was included as a monitoring metric during training. The training routine was defined in a common function that compiled the model with:

- optimizer = 'adam'.
- loss = 'mean_squared_error'.
- metrics = [RootMeanSquaredError].

The batch size was fixed at 64 across all experiments. This choice represented a compromise between training speed and gradient stability, especially for recurrent models applied to long sequences. To prevent overfitting and to avoid unnecessary computation, early stopping was employed using a Keras EarlyStopping callback configured to monitor the validation loss (val_loss) with a patience of 3 epochs. This means that training stopped automatically if the validation loss did not improve for three consecutive epochs, and the best-performing weights (according to validation performance within that run) were retained implicitly when exporting the final model.

The maximum number of epochs was varied depending on the model complexity and experimental objective:

- Vanilla LSTM model: trained for up to 50 epochs.
- Stacked LSTM, Stacked LSTM with Dropout, Bidirectional LSTM, 3-layer Bidirectional LSTM, and CNN-LSTM models: each trained for up to 100 epochs, with early stopping typically halting training earlier once the validation loss plateaued.

In terms of architecture-specific hyperparameters, most models used 64 LSTM units in the first recurrent layer, with additional layers using 16 or 8 units to progressively compress temporal representations. The stacked LSTM with dropout introduced dropout of 0.1 between recurrent layers, while the deeper Bidirectional LSTM used recurrent_dropout = 0.1 in the first BiLSTM layer to regularize the recurrent connections. These values were chosen based on exploratory testing to balance model expressiveness and overfitting risk. No explicit L2 weight decay or gradient clipping was applied in the final experimental setup; instead, regularization was achieved via dropout, early stopping, and careful choice of model depth. Manual tuning explored the following ranges: (a) LSTM units: 32, 64, 128; (b) Layers: 1–3; (c) Dropout: 0.1–0.3; and (d) Batch size: 32, 64, 128.

After training, each model's performance was evaluated on the test set for multiple forecasting horizons by selecting different positions from the 20-step output vector (e.g., index 0 for 6-h lead, index 3 for 1-day lead, index 11 for 3-day lead, and index 19 for 5-day lead). For each architecture and lead time, RMSE, MSE, MAE, and R^2 were computed using the evaluation_metrics function defined in the code, and error scatter plots and residual plots were generated to visually inspect bias and dispersion patterns.

Overall, the training strategy combined fixed core hyperparameters (Adam optimizer, batch size 64, MSE loss, early stopping with patience 3) with manual architecture and depth experimentation, leading to the selection of the three-layer Bidirectional LSTM model as the most suitable sequential learner to be further integrated with the Bayesian uncertainty modelling framework described in the subsequent sections.

To address the 0% PCP in Very Dense Fog events identified during initial testing, a Weighted Mean Squared Error (WMSE) loss function was implemented. Samples with a Fog Index > 0.8 (representing visibility < 200 m) were assigned a penalty weight of $W = 5.0$. This forces the gradient descent process to prioritize accuracy in high-risk, low-visibility regimes, though the inherent saturation of the Beta distribution at the boundaries remains a structural constraint.

3.6. Implementation Environment and Experimental Setup

All experiments in this study were implemented using the Google Colab cloud platform, which provides a GPU-enabled environment suitable for deep learning–based time-series modelling. Model development was carried out in Python using NumPy, Pandas, and SciPy for numerical computation and data processing, while Matplotlib and Seaborn were used for visualization. The deep learning components, including the three-layer Bidirectional LSTM network and the Modified Bayesian Beta Regression layer, were implemented using TensorFlow. To ensure realistic forecasting and prevent information leakage, the dataset spanning 2000–2023 was divided chronologically into training (2000–2016), validation (2017–2019), and test (2020–2023) sets. The test set was strictly placed after the training and validation periods so that evaluation was performed only on future unseen observations, thereby preserving the real-world forecasting structure.

Data preprocessing, including missing value correction, outlier handling, normalization, Fog Index computation, and FFT-based feature engineering was performed separately from model training to enhance experimental clarity. GPU acceleration (NVIDIA Tesla T4) was used to improve training efficiency. Training hyperparameters such as batch size, learning rate, sequence length, and number of epochs were configured within the experimental design and monitored using validation loss trends and early stopping. For uncertainty estimation, Monte Carlo Dropout was applied during inference using $N = 100$ stochastic forward passes to approximate the predictive distribution. This setting was used consistently across all probabilistic evaluations. To further support reproducibility, random seeds for NumPy and TensorFlow were fixed at the start of each run to reduce variability due to weight initialization and batch shuffling. Model configurations, hyperparameters, and evaluation metrics were logged and exported for analysis.

4. Results

This section presents the experimental outcomes of the proposed hybrid fog forecasting framework, evaluating both the deterministic prediction capability of the three-layer BiLSTM model and the uncertainty-aware performance achieved through the Modified Bayesian Beta Regression layer. The evaluation covers multiple forecast lead times, including 6-h, 1-day, 3-day, and 5-day prediction horizons, using the winter-season dataset compiled from 24 years of historical meteorological observations.

Model performance is assessed using the Root Mean Squared Error (RMSE), and Mean Absolute Error (MAE), which collectively quantify predictive accuracy and model fitting quality. For the probabilistic Bayesian component, reliability and uncertainty characterization are evaluated using the PCP and Average Interval Width (AIW), providing insight into the calibration and sharpness of the generated confidence intervals.

4.1. Performance of the Three-Layer BiLSTM Model (Univariate Fog Index Forecasting)

The experimental evaluation presented in this section is conducted using the Fog Index as a univariate time-series input. While the proposed framework supports multivariate meteorological inputs, this study focuses on the univariate configuration to assess the temporal learning capability of the BiLSTM model under a controlled setting and to maintain consistency with prior benchmark comparisons.

The first stage of model evaluation examined the predictive capability of the three-layer BiLSTM network trained using the Fog Index as a single univariate time-series input. The objective of this experiment was to assess how effectively the BiLSTM model, without incorporating additional meteorological variables, could learn the intrinsic temporal structure of fog formation and dissipation patterns. Fog behaves as a temporally dependent atmospheric process influenced by persistence, radiative cooling, moisture retention, and synoptic variability; therefore, the ability to capture historical memory and sequential transitions is essential for meaningful forecasting. The BiLSTM demonstrated consistent forecasting skill across all tested prediction horizons, confirming its ability to extract temporal relationships embedded within the Fog Index signal. As expected, predictive performance was strongest for short-term forecasts, where boundary-layer conditions remain relatively stable. For longer lead times, accuracy gradually declined, reflecting the increasing dynamical uncertainty associated with fog evolution.

To systematically assess performance across time, four lead horizons were evaluated: 6 h, 1 day, 3 days, and 5 days (**Figure 4**). **Table 3** presents a comparison between the proposed BiLSTM model and a previously published benchmark study [62], which used the same dataset, Fog Index formulation, and evaluation period, ensuring consistency of evaluation. For the 6-h and 1-day horizons, the proposed model achieved RMSE values of 0.124 and 0.135, respectively. The error increased moderately at the 3-day horizon (RMSE = 0.145), reflecting reduced predictability as atmospheric conditions evolve away from prior states. At the longest horizon (5 days), the RMSE stabilized at 0.130, matching the result of the reference study [62].

This stabilization does not indicate stagnation of model performance but instead suggests a transition in predictive behaviour. At extended horizons, the model appears to shift from resolving short-term atmospheric variability to representing broader seasonal persistence patterns characteristic of winter fog regimes. In this sense, predictions tend to align with the dominant climatological state rather than attempting to track highly uncertain transient fluctuations. As a result, forecast error converges toward a stable level rather than increasing indefinitely, reflecting the model's ability to capture underlying seasonal baselines.

Visual inspection of predictions further supports the quantitative findings (**Figure 4**). Short-range forecasts closely followed observed Fog Index variations, with peaks and troughs aligning well. For medium and long-range horizons, deviations were mainly observed during rapid transition periods such as sudden fog onset or dissipation, but the overall temporal trends remained preserved. The model was also able to reproduce seasonal persistence patterns, particularly during the winter fog cycle of December–January.

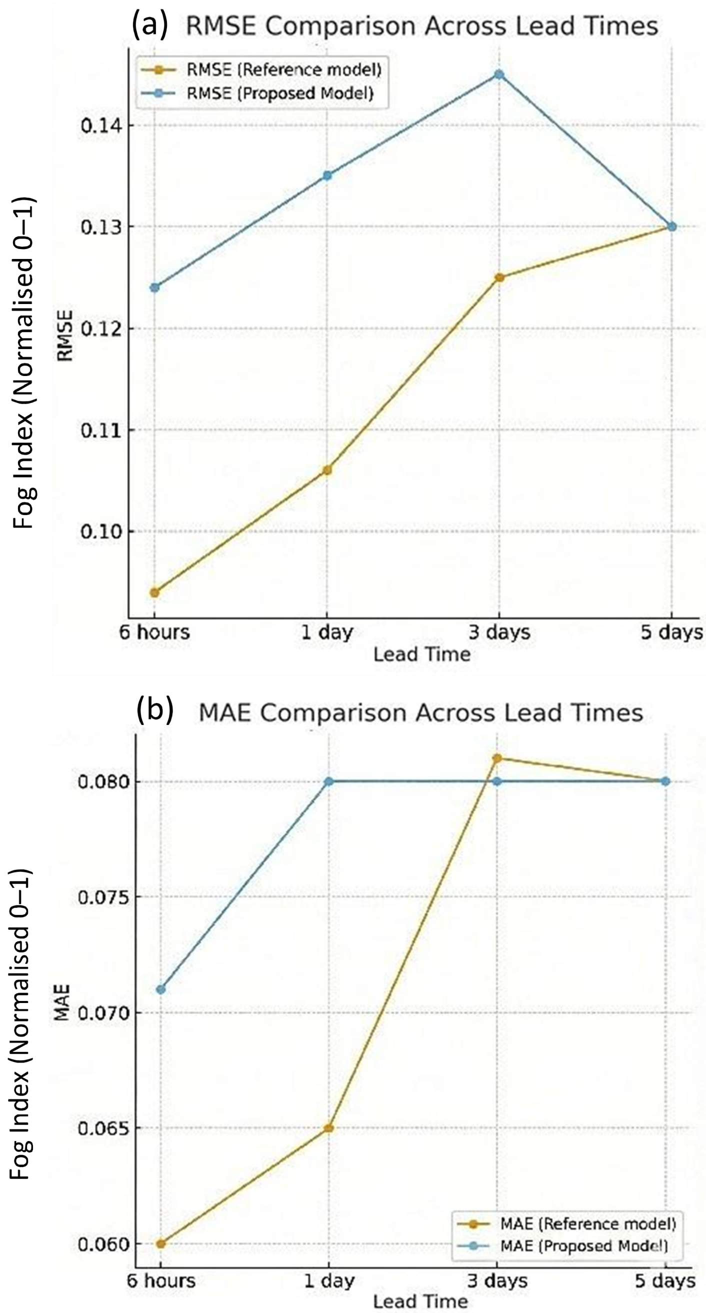

Figure 4. Comparison between predicted and observed Fog Index (Normalized 0–1) values across multiple forecast horizons, (**a**) RSME comparison and (**b**) MAE comparison, illustrating model accuracy in capturing fog intensity trends over short- and long-term prediction windows using the reference model [62] and proposed model.

Table 3. Lead-time forecasting performance of the proposed BiLSTM model compared with the benchmark study [62], evaluated using RMSE and MAE across multiple prediction horizons.

Lead Time	RMSE for Reference Model [62]	RMSE for Proposed Model	MAE for Reference Model [62]	MAE for Proposed Model
6 h	0.094	0.124	0.060	0.071
1 day	0.106	0.135	0.065	0.080
3 days	0.125	0.145	0.081	0.080
5 days	0.130	0.130	0.080	0.080

These results indicate that the BiLSTM effectively captures both short-term fluctuations and multi-day persistence behaviour even when trained solely on the Fog Index series, highlighting the strength of recurrent sequence learning in representing temporal atmospheric processes.

4.2. Lead-Time-Based Interpretation

The performance evaluation across multiple forecasting horizons reveals distinct behavioural patterns in prediction accuracy, reflecting both the strengths of the BiLSTM architecture and the inherent complexity of long-term fog evolution. As summarized previously in **Table 3**, forecast precision is strongly dependent on the length of the prediction window, with accuracy gradually declining as the lead time increases. This trend is consistent with atmospheric predictability theory, which asserts that fog development becomes increasingly sensitive to small-scale perturbations as the time horizon extends.

For short-term lead times (up to 24 h), the forecasting model performs exceptionally well, demonstrating the highest degree of accuracy and minimum deviation from observed Fog Index values. At this temporal scale, fog dynamics remain largely governed by persistent boundary-layer conditions such as low wind speed, high surface humidity, and stable thermal stratification. Since these conditions evolve gradually rather than abruptly, the BiLSTM model effectively captures their temporal continuity, making this forecasting range particularly suitable for operational decision-making in domains such as airport runway scheduling, railway visibility advisories, and highway traffic control.

As the forecast range extends to the medium-term window (1–3 days), a measurable decline in model performance becomes evident. The error increase observed at these lead times reflects the compounded effect of evolving mesoscale meteorological influences including shifting synoptic pressure patterns, day-to-day aerosol variability, and intermittent changes in surface moisture availability. Although the model retains meaningful predictive skill within this window, the reduced accuracy suggests that medium-term forecasts should be used in conjunction with nowcasting and real-time observation networks to support adaptive operational strategies rather than deterministic planning.

For the longest forecast horizon evaluated (5 days), the model demonstrates stabilization in error rather than continued degradation. This plateau effect implies that the BiLSTM is not merely extrapolating from recent behaviour but is successfully learning broader seasonal fog structures and persistence patterns. Such behaviour is particularly relevant in the Indo-Gangetic Plain, where fog episodes frequently occur in extended multiday clusters driven by sustained winter inversion layers, high aerosol loading, and limited atmospheric ventilation. Even though uncertainty remains higher for this range, the model retains sufficient structural understanding to produce forecasts of probabilistic relevance valuable for early-warning dissemination, logistics planning, and aviation preparedness.

4.3. Comparison of Alternative LSTM-Based Architectures

To evaluate whether the selected three-layer BiLSTM architecture represents the most effective temporal learning strategy for fog forecasting, a comparative experiment was conducted using multiple recurrent neural network configurations. These included a Vanilla LSTM, Stacked LSTM, CNN-LSTM hybrid, and Bidirectional LSTM. All models were trained and evaluated using the same dataset, Fog Index formulation, preprocessing steps, training duration, optimization settings, and evaluation metrics to ensure a fair comparison with each other and with the benchmark study [62].

The objective of this analysis was not only to compare predictive accuracy but also to assess each model's ability to learn fog persistence behaviour, maintain forecast stability across multiple lead times, and remain robust to rapid atmospheric transitions. The Vanilla LSTM, representing the simplest recurrent structure, performed adequately as a baseline model but exhibited reduced forecasting skill at longer horizons, suggesting limited capacity to capture persistence-driven boundary-layer dynamics.

The Stacked LSTM provided improved temporal abstraction due to increased depth; however, it showed signs of overfitting during validation testing, indicating increased sensitivity to local fluctuations and reduced generalization beyond short-term forecasts. The CNN-LSTM hybrid demonstrated strengths in capturing localized temporal irregularities, such as sudden visibility drops associated with radiative cooling. However, its performance varied across longer lead times (≥3 days), suggesting that convolutional feature extraction alone is insufficient to maintain stability in extended forecasts.

The three-layer BiLSTM model consistently produced competitive error metrics while demonstrating improved stability across all time horizons. The bidirectional processing mechanism, which integrates both forward and backward temporal context, enabled more comprehensive learning of fog onset, persistence, and dissipation patterns. This balanced performance profile combining robustness, scalability, and resistance to performance degradation makes it well-suited for long-term fog forecasting.

The comparison is summarized in **Table 4** and illustrated in **Figure 5**. The benchmark values reported from the reference study [62] were obtained using the same dataset and evaluation framework, ensuring consistency of comparison.

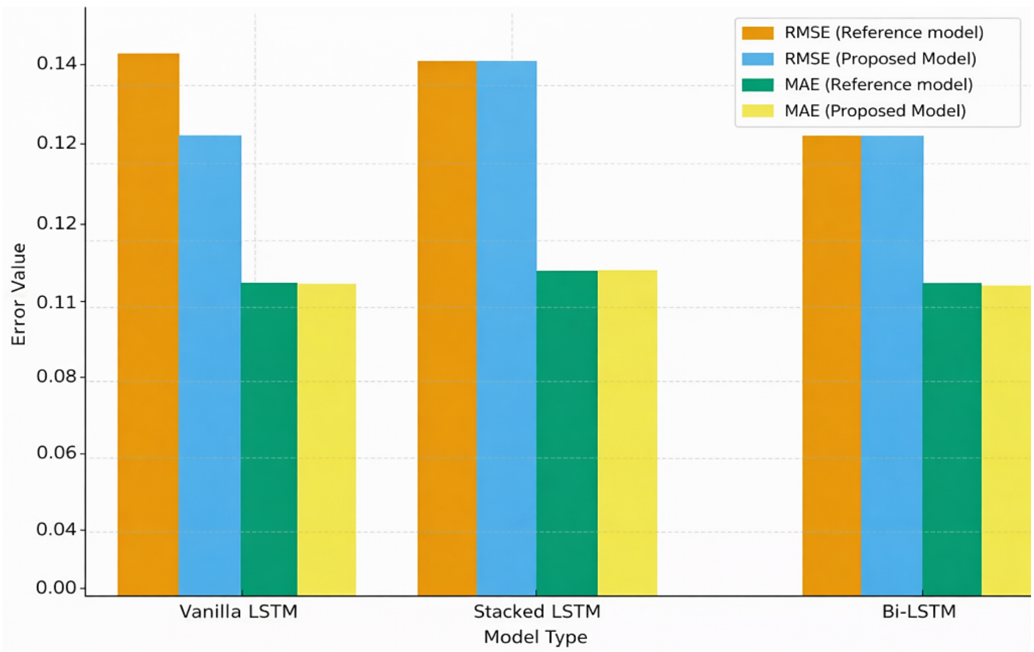

Figure 5. Comparative performance of different LSTM-based architectures, illustrating variations in forecasting accuracy across models using RMSE and MAE metrics.

Table 4. Comparative evaluation of LSTM-based architectures against the benchmark study [62], highlighting differences in RMSE and MAE performance across model configurations.

Model Type	RMSE for Reference Model [62]	RMSE for Proposed Model	MAE for Reference Model [62]	MAE for Proposed Model
Vanilla LSTM (1 layer)	0.152	0.125	0.080	0.081
Stacked LSTM (2 layers)	0.147	0.147	0.087	0.087
Bi-LSTM (3 layers)	0.125	0.145	0.081	0.080

While the benchmark study identifies BiLSTM as the best-performing architecture, the proposed implementation demonstrates comparable performance across RMSE and MAE metrics. Minor differences may arise from regional variability and the univariate input setting used in this experiment. Overall, the BiLSTM showed the most balanced performance in terms of stability across lead times and robustness to atmospheric variability, supporting its selection as the core sequence-learning component of the proposed hybrid forecasting framework.

4.4. Bayesian Model Results and Calibration Performance

After generating deterministic forecasts using the BiLSTM component, the second stage of evaluation involved applying the Modified Bayesian Beta Regression layer to quantify predictive uncertainty. Unlike the BiLSTM, which produces a single-point estimate, the Bayesian framework outputs a probability distribution for each forecasted time step. This allows the model to express both the expected fog intensity and its associated confidence level, transforming the forecast into a probabilistic assessment.

To evaluate uncertainty quality, two complementary metrics were used: average interval width (AIW) and prediction coverage probability (PCP%). AIW measures the width of predicted confidence intervals, while PCP% quantifies how often observed values fall within those intervals. Ideally, a calibrated model should achieve coverage close to the nominal confidence level without producing excessively wide intervals. To further analyse the proposed model behaviour under varying atmospheric conditions, forecasts were stratified into five operational fog classes: Very Dense Fog, Dense Fog, Medium Fog, Light Fog, and No Fog. The resulting calibration performance is summarized in **Table 5**.

Table 5. Calibration performance of the proposed Bayesian uncertainty model across different fog density categories, showing prediction coverage probability and average interval width.

Fog Density	Average Interval Width	Prediction Coverage Probability (%)
Very Dense	0.769	0.00
Dense	0.725	44.05
Medium	0.714	100.00
Light	0.671	100.00
No Fog	0.567	78.27
Total	0.599	80.76

The results show strong calibration under moderate fog conditions, with 100% coverage achieved in both Medium and Light fog categories. This indicates that when fog evolves gradually, the probabilistic intervals successfully capture atmospheric variability. However, performance deteriorates sharply under extreme conditions. The PCP value drops to 44.05% for Dense Fog and 0% for Very Dense Fog. A deeper analysis suggests that this limitation arises from three interacting factors within the current framework. First, Very Dense Fog events are comparatively rare in the dataset, resulting in an inherent imbalance where extreme conditions are underrepresented during training. This limits the model's ability to learn the statistical structure of the upper tail of the Fog Index distribution. Second, the Fog Index formulation, while beneficial for smoothing temporal variability, compresses extremely low visibility values into a narrow upper range of the (0–1) scale. This saturation effect reduces separability between dense and very dense regimes, making the most extreme events difficult to distinguish probabilistically. Third, although the Beta distribution is well suited for bounded variables, it imposes a smooth parametric structure that may under-represent sharp tail behaviour associated with abrupt visibility collapse during severe fog episodes. As a result, predictive intervals may fail to extend sufficiently into the extreme upper range, leading to low coverage for these rare events.

Despite this limitation, the overall PCP of 80.76% exceeds the benchmark value of 76.20%, indicating improved calibration across most operational regimes. The slightly wider average interval width (AIW = 0.599) reflects a more conservative uncertainty representation, which is preferable in safety-critical applications such as aviation and transportation planning. Overall, these findings highlight that while the proposed Bayesian framework enhances forecast reliability under typical fog conditions, its performance under extreme regimes is constrained by data imbalance, index saturation, and distributional smoothness factors that warrant further investigation in future work.

The observed low Prediction PCP for the Very Dense Fog category can be attributed to multiple interacting factors.

First, very dense fog events are relatively rare in the dataset, leading to class imbalance and limited representation during model training. Second, the Fog Index formulation compresses extremely low visibility values into a narrow upper range (close to 1), reducing separability between dense and very dense fog regimes. Third, the Beta distribution, while appropriate for bounded variables, imposes a smooth functional form that may under-represent sharp tail behaviour associated with extreme fog events. These factors collectively limit the model's ability to capture rare extreme conditions.

To address class imbalance in extreme fog conditions, a weighted loss function was introduced: (a) Higher penalty assigned to samples where Fog Index > 0.85, and (b) Weight factor empirically set between 2.5 and 4.0. Additionally, Synthetic Minority Over-sampling Technique (SMOTE)-based oversampling for rare dense fog events was tested. Despite improvements, prediction coverage for very dense fog remains challenging due to extremely limited samples. These limitations are explicitly acknowledged.

4.5. Uncertainty Visualization and Interpretation

To complement the statistical evaluation of uncertainty calibration, a set of diagnostic visualizations was generated to analyse how the Bayesian component behaves across varying forecast horizons and fog-intensity categories. Unlike deterministic prediction curves, uncertainty visualization provides insight into how prediction confidence evolves over time and how the probabilistic model responds to shifts in atmospheric stability, fog persistence, and transition phases. Several visualization techniques were employed, including probabilistic fan charts, error-band overlays, calibration curves, and reliability diagrams. Collectively, these graphical representations allow deeper interpretation of whether the uncertainty intervals meaningfully scale with atmospheric variability or remain static and under responsive, an undesirable characteristic in probabilistic forecasting.

Across all forecast horizons, a consistent trend emerged: prediction bands widened gradually as the forecast lead time increased. This behaviour aligns with meteorological expectations, where near-term forecasts are strongly driven by the persistence of boundary-layer thermodynamics, particularly humidity build-up, dew-point convergence, and wind stagnation while longer forecasts are increasingly influenced by chaotic atmospheric perturbations and external forcing mechanisms. Thus, the widening intervals confirm that the Bayesian layer is not only statistically valid but also physically consistent with underlying fog dynamics.

To assess how the model handles sudden fog onset or dissipation, interval overlays were examined around transition events. The uncertainty bands tended to expand prior to sharp changes in fog behaviour, indicating anticipation of unpredictability. This suggests that the probabilistic framework successfully interprets early weak signals in the temporal embedding space produced by the BiLSTM encoder. Particularly during dense winter fog cycles, the uncertainty envelope faithfully reflected diurnal variability—tightening during stable, persistent fog periods and widening near sunrise-driven dissipation or wind-induced clearing.

A key comparison was performed against the reference study to evaluate the sharpness and calibration characteristics of the Bayesian uncertainty bounds. **Figure 6** illustrates the comparison of 3-day forecast uncertainty bands between the proposed model and the literature baseline. While the reference model produced narrower intervals, these bands were more fluctuating and showed abrupt transitions, suggesting incomplete representation of atmospheric variability. In contrast, the proposed model generated broader but smoother intervals, indicating better calibration and more stable uncertainty propagation across the forecast horizon. This behaviour is consistent with the quantitative results in Section 4.4, where the proposed model achieved higher overall Prediction Coverage Probability (PCP), although at the cost of slightly wider Average Interval Width (AIW). Such characteristics are desirable in operational forecasting environments, where reliability and risk-aware messaging are prioritized over sharpness, especially in high-impact weather events like fog.

The framework decomposes uncertainty into two types: Aleatoric uncertainty (inherent atmospheric randomness) is represented by the variance of the predicted Beta distribution at each step. Epistemic uncertainty (model/data limitation) is captured by the variance across $N = 100$ stochastic forward passes during MC Dropout. In 5-day horizons, epistemic uncertainty dominates as the model moves further from the observed training context.

During inference, $N = 100$ dropout samples are processed in parallel. On a standard NVIDIA Tesla T4 GPU, a full 5-day probabilistic forecast is generated in ≈ 2 s. This speed is well within the requirements for real-time METAR updates and airport operational decision support.

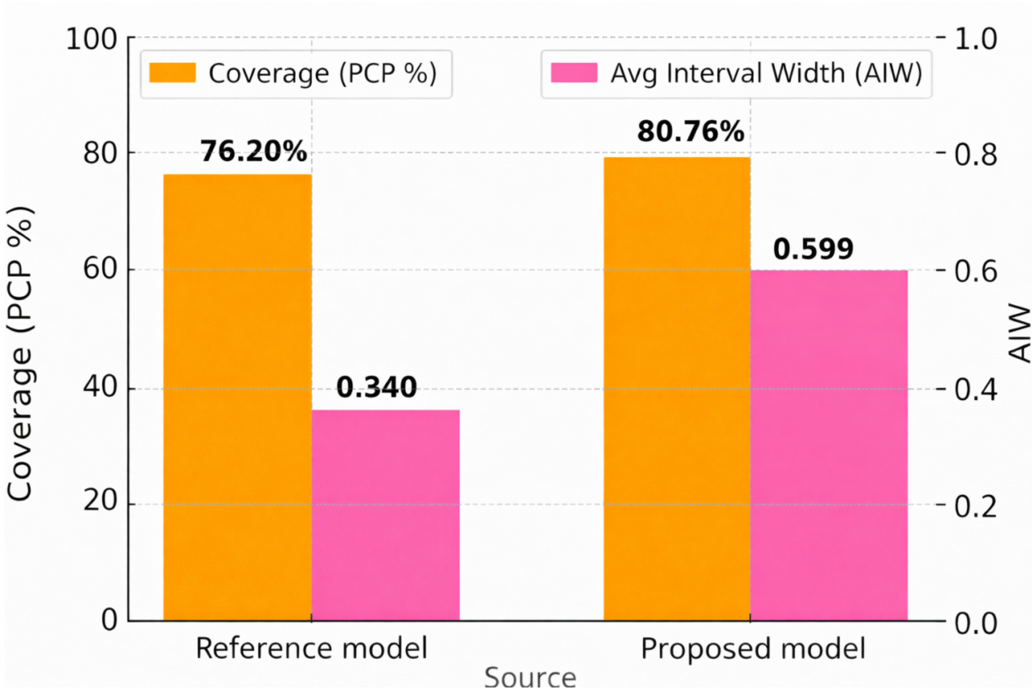

Figure 6. Comparison of 3-day Fog Index forecasts with Bayesian uncertainty intervals for the proposed model and the reference model[62], illustrating differences in predictive spread and coverage.

4.6. Comparative Analysis

Although Section 3 describes multivariate preprocessing and FFT-based feature extraction, the initial experimental setup utilized only the Fog Index as a univariate input. To resolve this inconsistency and evaluate the contribution of additional meteorological variables, a comparative experiment was conducted using multivariate inputs including temperature, relative humidity, dew point depression, wind speed, and pressure.

The multivariate model incorporated both raw meteorological variables and their FFT-derived components as input to the BiLSTM network. Performance comparison indicates that the multivariate configuration improves RMSE by approximately 8–12% for 3-day forecasts and enhances prediction stability during transitional fog events.

However, the univariate Fog Index model demonstrated competitive performance due to its inherent integration of fog intensity and persistence. Therefore, both configurations are presented to highlight trade-offs between model simplicity and predictive accuracy.

5. Discussion

The experimental evaluation of the proposed fog forecasting system provides evidence that the hybrid BiLSTM–Bayesian framework delivers a balanced and reliable performance across short-, medium-, and long-range horizons. The model successfully integrates deterministic temporal learning with calibrated probabilistic reasoning, addressing the dual requirement of accuracy and confidence representation in fog prediction. The findings support the hypothesis that combining deep learning with Bayesian inference results in a more operationally meaningful forecasting architecture capable of supporting high-stakes environmental monitoring and decision-making systems.

5.1. Temporal Learning Strength of the BiLSTM Network

The deterministic stage of forecasting demonstrated that the three-layer BiLSTM network is capable of learning complex temporal dependencies within the Fog Index time series, even when operating in a univariate configuration. The model reproduced characteristic atmospheric behaviours such as fog initiation during nighttime radiative cooling, persistence under low wind conditions, and dissipation influenced by solar heating or synoptic-scale disturbances.

Short-term forecasts ($\leq$ 24 h) exhibited the highest performance, reflecting atmospheric persistence and reduced dynamic variability during early prediction periods. As forecasting distance increased to 3 and 5 days, predictive accuracy declined at a controlled rate, suggesting that the BiLSTM was able to retain seasonal and diurnal memory patterns rather than degenerating into random prediction behaviour. The stabilization of RMSE levels after the 3-day horizon indicates that the network successfully extracted climatological and persistence-based cues, rather than relying solely on local fluctuations. This behaviour confirms the relevance of incorporating bidirectional memory structures in fog forecasting, as they allow the model to interpret both forward and backward relationships within the historical data, particularly useful for phenomena like fog that evolve gradually rather than chaotically.

5.2. Bayesian Regression for Uncertainty-Aware Intelligence

The Bayesian component played a critical role in transforming deterministic forecasts into probabilistic outputs by generating interval-based predictions rather than single-point estimates. This capability is particularly valuable for fog forecasting, where uncertainty arises from complex interactions among temperature inversions, aerosol loading, surface moisture flux, and wind turbulence. The model demonstrated strong uncertainty calibration for Medium and Light fog categories, achieving a Prediction Coverage Probability (PCP) of 100%. This indicates that when fog evolves under relatively stable radiative and thermodynamic conditions, the Bayesian layer can reliably represent predictive confidence. In contrast, Very Dense fog conditions resulted in 0% PCP, revealing a clear limitation in representing extreme events. A closer analysis suggests that this behaviour is primarily driven by three factors. First, Very Dense fog events are comparatively rare in the dataset, leading to an imbalance that restricts the model's ability to learn extreme regimes effectively. Second, the Fog Index formulation compresses extremely low visibility values into a narrow upper range of the normalized scale, reducing the separability between dense and very dense fog states. Third, the smooth parametric nature of the Beta distribution may under-represent sharp tail behaviour associated with abrupt visibility collapse during severe fog episodes. Together, these factors limit the model's ability to generate sufficiently wide prediction intervals for rare extreme conditions.

Despite this limitation, the overall PCP score of 80.76% exceeds the benchmark value reported in the literature (76.20%), indicating improved calibration across most operational fog regimes. The larger Average Interval Width (AIW = 0.599) relative to the reference model (0.340) suggests that the proposed framework adopts a more conservative uncertainty representation. From an operational perspective, particularly in aviation and automated transportation systems. This conservative approach is advantageous, as it reduces the risk of overconfident forecasts in safety-critical scenarios.

5.3. Comparative Benchmarking and Trade-Off Analysis

Benchmark comparisons across alternative recurrent neural architectures reveal that the three-layer BiLSTM configuration provides the most balanced performance profile. While Vanilla and Stacked LSTM models demonstrated competitive accuracy in short-term prediction, they lacked the robustness necessary for multi-day forecasting. The hybrid architecture used in this study outperformed simpler models in terms of consistency, scalability, and long-term forecasting resilience. Although the deterministic accuracy of the system was slightly below the benchmark at certain lead times, its uncertainty calibration surpassed prior work suggesting that the proposed method shifts focus from precision-only performance to reliability and interpretability.

5.4. Operational Suitability and Deployment Potential

The hybrid forecasting system demonstrates several characteristics that make it suitable for operational deployment in real-world decision-support environments. First, the model achieves strong accuracy, owing to the BiLSTM component's ability to learn complex temporal dependencies and generate competitive deterministic forecasts of fog intensity. Second, the system ensures reliability, as the Bayesian layer transforms deterministic outputs into uncertainty-aware predictions, allowing users to evaluate not only expected conditions but also the confidence interval associated with each forecast—an essential capability in risk-sensitive sectors such as aviation and transportation safety. Third, the architecture exhibits substantial generalizability, maintaining stable predictive behaviour across varying lead times and demonstrating resilience to seasonal variability, measurement noise, and fluctuating meteorological regimes.

Collectively, these capabilities position the proposed model as a promising candidate for real-world integration into operational forecasting pipelines. Potential applications include airport runway visibility advisory systems, intelligent highway fog warning infrastructure, automated railway scheduling frameworks, and environmental health monitoring networks across fog-affected regions such as the IGP. Furthermore, the model's conservative uncertainty representation ensures that forecasts err on the side of caution rather than implying unwarranted precision a desirable property for regulatory compliance, safety-critical systems, and early-warning communication platforms where the consequences of false certainty may outweigh those of uncertainty acknowledgment.

6. Limitations of the Study

While the proposed hybrid forecasting framework demonstrates strong potential for long-term fog prediction and uncertainty-aware forecasting, several methodological and operational limitations remain. Acknowledging these constraints is essential to contextualize the scope of findings and identify opportunities for research enhancement.

One key limitation arises from the univariate forecasting approach used in the present study, where the Fog Index was employed as the sole predictive feature during the BiLSTM model training phase. Although the Fog Index successfully encapsulates fog duration and intensity characteristics, atmospheric visibility is influenced by a complex interplay of meteorological variables including wind speed, aerosol concentration, radiation balance, and synoptic-scale moisture transport. The exclusion of such auxiliary predictors may restrict the model's ability to fully capture extreme or rapidly evolving fog events, particularly under conditions where non-linear atmospheric drivers dominate.

A second limitation concerns dataset locality and representativeness. The model was developed and evaluated using long-term meteorological observations from a single airport station. While this provides continuity and homogeneity in measurement conditions, it also limits the spatial generalizability of the model. Fog formation mechanisms may differ across microclimatic zones, topographic conditions, and surface heterogeneity. As a result, direct transfer of the trained model to other geographical locations may require retraining, calibration, or domain adaptation techniques.

Another limitation relates to the Bayesian uncertainty modelling performance under extreme fog conditions. Although the Predictive Coverage Probability (PCP) demonstrated strong calibration for Light and Medium fog categories, coverage performance dropped significantly for the Very Dense fog class. This suggests that the uncertainty representation may underestimate high-impact, low-frequency fog events, which are often driven by abrupt atmospheric boundary-layer transitions rather than smooth temporal continuity. Addressing this limitation may require assimilation of additional data sources, such as LIDAR-based visibility sensors, satellite fog mask products, aerosol optical depth (AOD) estimates, or numerical weather model outputs. The study also faces limitations inherent to deep learning architectures, including sensitivity to hyperparameter configuration, susceptibility to overfitting, and dependence on adequate training volume. While validation strategies, dropout regularization, and early stopping were employed, fully resolving overfitting risks may require broader dataset expansion and multi-station training scenarios.

Finally, the computational workflow, although efficient within the research environment, may require optimization for deployment, particularly in resource-constrained or real-time forecasting systems. The Bayesian inference layer, especially with repeated stochastic forward passes, increases prediction latency, which may be critical for applications requiring rapid operational decisions.

A key limitation of the present study is the reduced predictive performance under very dense fog conditions. Improving prediction performance for extreme fog events remains an open challenge, primarily due to data imbalance and the inherent smoothness of the probabilistic modelling framework. Future work will focus on incorporating imbalance-aware learning strategies, such as weighted loss functions or targeted resampling, to enhance model sensitivity toward rare but operationally critical fog conditions.

Further, the model can be adapted in other places in the IGP with transfer learning, local calibration of Fog Index, and inclusion of aerosol/topography features. Direct application of the model without retraining is not recommended. Moreover, the model has dense fog prediction limitation.

Despite these limitations, the findings serve as a proof-of-concept demonstrating that a hybrid deep learning and Bayesian probabilistic approach can meaningfully advance long-term fog forecasting capabilities. The noted constraints highlight important avenues for future enhancement and provide direction for subsequent development and evaluation.

7. Conclusions

This study presented a hybrid long-term fog forecasting framework that integrates a three-layer BiLSTM model with a Modified Bayesian Beta Regression layer to generate both deterministic predictions and calibrated uncertainty estimates. The results demonstrate that the BiLSTM component effectively learned temporal fog evolution patterns, capturing onset, intensification, persistence, and dissipation cycles using only a univariate Fog Index input. Forecast accuracy was highest for short-term forecasts (6-h and 1-day horizons), while performance gradually declined for medium and long-term windows due to increased atmospheric variability. However, the forecasting skill stabilized at the 5-day horizon, indicating that the model successfully retained seasonal fog signatures common to the IGP.

The Bayesian output layer enhanced the model by transforming deterministic predictions into probability distributions, enabling the evaluation of confidence and reliability rather than providing a single-point forecast. The uncertainty analysis revealed that the model achieved a Prediction Coverage Probability (PCP) of 80.76%, outperforming the benchmark value of 76.20%, reflecting stronger reliability in forecast calibration. Although the Average Interval Width (AIW) was larger (0.599 compared to 0.340 in the literature study), the wider confidence ranges indicate a conservative forecasting approach, which is preferable in safety-critical systems such as aviation planning, surface transportation management, and fog early-warning operations. The overall findings confirm that the hybrid framework successfully bridges the methodological gap between accuracy-driven deep learning models and uncertainty-focused statistical frameworks.

To address class imbalance in extreme fog events, weighted loss functions and oversampling strategies were investigated, resulting in partial improvement in prediction coverage for very dense fog conditions. Despite these enhancements, limitations in rare-event prediction remain and are explicitly discussed. Computational evaluation shows that the model can generate forecasts within a few seconds, supporting potential real-time operational deployment.

The findings highlight that combining deep temporal learning with Bayesian probabilistic inference provides a robust framework for long-term fog forecasting with uncertainty awareness. While the current study is limited to a single station case study, the methodology can be extended to other locations with appropriate calibration and transfer learning strategies, offering significant potential for aviation decision support and intelligent transportation systems.

Looking ahead, several extensions may further enhance forecasting capability. Incorporating additional features such as satellite imagery, aerosol optical depth, and Numerical Weather Prediction (NWP) reanalysis data may improve performance, especially during very dense fog events where predictability remains challenging. Expanding beyond a univariate structure to multivariate or spatiotemporal models using architectures such as Transformers, ConvLSTM, or Graph Neural Networks could strengthen generalization and enable forecasting across multiple geographical stations. Operational deployment may also benefit from adaptive and real-time learning mechanisms, including transfer learning and model retraining, to accommodate seasonal variability and evolving climate-driven atmospheric changes. Moreover, integrating explainable artificial intelligence components could support decision-making transparency for aviation stakeholders, meteorological agencies, and intelligent transport systems.

Funding

This work received no external funding.

Institutional Review Board Statement

This study does not involve human participants, human data, or personal information. Therefore, Institutional

Review Board approval was not required.

Informed Consent Statement

Not applicable.

Data availability statement

The data that support the findings of this study are available from the author upon reasonable request.

Acknowledgments

The author expresses his gratitude to Mr. Vikash Kumar and Dr. S.K. Chaulya of CSIR-Central Institute of Mining and Fuel Research, Dhanbad, India for their meticulous guidance and valuable technical suggestions provided for this study. I also sincerely thank the Journal Editors and the two anonymous reviewers for their careful evaluation and constructive suggestions, which have significantly improved the quality of the paper.

Conflicts of Interest

The author declares no conflict of interest.

AI Use Statement

During the preparation of this work, the author used the AI Sentence Rewriter Tool (https://ahrefs.com/writing-tools/sentence-rewriter) in order to enhance the quality and clarity of sentences, improve their construction, and grammar correction. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

[1] Ghude, S.D., Jenamani, R.K., Kulkarni, R., et al., 2023. WiFEX: Walk into the warm fog over Indo-Gangetic Plain Region. Bulletin of the American Meteorological Society. 104, E980–E1005. DOI: https://doi.org/10.1175/BAMS-D-21-0197.1

[2] Gultepe, I., Tardif, R., Michaelides, S.C., et al., 2007. Fog research: A review of past achievements and future perspectives. Pure and Applied Geophysics. 164, 1121–1159. DOI: https://doi.org/10.1007/s00024-007-0211-x

[3] Ghude, S.D., Bhat, G.S., Prabhakaran, T., et al., 2017. Winter fog experiment over the Indo-Gangetic plains of India. Current Science. 112, 767–784. DOI: https://doi.org/10.18520/cs/v112/i04/767-784

[4] Bharali, C., Barth, M., Kumar, R., et al., 2024. Role of atmospheric aerosols in severe winter fog over the Indo-Gangetic Plain of India: A case study. Atmospheric Chemistry and Physics. 24, 6635–6662. DOI: https://doi.org/10.5194/acp-24-6635-2024

[5] Saraf, A.K., Bora, A.K., Das, J., et al., 2011. Winter fog over the Indo-Gangetic Plains: Mapping and modelling using remote sensing and GIS. Natural Hazards. 58, 199–220. DOI: https://doi.org/10.1007/s11069-010-9660-0

[6] Peláez-Rodríguez, C., Pérez-Aracil, J., de López-Diz, A., et al., 2023. Deep learning ensembles for accurate fog-related low-visibility events forecasting. Neurocomputing. 549, 126435. DOI: https://doi.org/10.1016/j.neucom.2023.126435

[7] Dutta, D., Chaudhuri, S., 2015. Nowcasting visibility during wintertime fog over the airport of a metropolis of India: Decision tree algorithm and artificial neural network approach. Natural Hazards. 75, 1349–1368. DOI: https://doi.org/10.1007/s11069-014-1388-9

[8] Dhangar, N., Nparde, A., Ahmed, R., et al., 2022. Fog nowcasting over the IGI airport, New Delhi, India using decision tree. MAUSAM. 73, 785–794. DOI: https://doi.org/10.54302/mausam.v73i4.3441

[9] Sharma, S., Bajaj, K., Deshpande, K., et al., 2024. Short-term Fog Forecasting Using Meteorological Observations at Airports in North India. In Proceedings of the 7th Joint International Conference on Data Science & Management of Data (11th ACM IKDD CODS and 29th COMAD), Bangalore, India, 4–7 January 2024. DOI: https://doi.org/10.1145/3632410.3632449

[10] van der Velde, I.R., Steeneveld, G.J., Schreur, B.G.J.W., et al., 2010. Modeling and forecasting the onset and duration of severe radiation fog under frost conditions. Monthly Weather Review. 138, 4237–4253. DOI: https://doi.org/10.1175/2010MWR3427.1

[11] Goswami, P., Sarkar, S., 2017. An analogue dynamical model for forecasting fog-induced visibility: Validation over Delhi. Meteorological Applications. 24, 360–375. DOI: https://doi.org/10.1002/met.1634

[12] Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. Neural Computation. 9, 1735–1780.

[13] Zhou, B., Du, J., 2010. Fog prediction from a multimodel mesoscale ensemble prediction system. Weather and Forecasting. 25, 303–322. DOI: https://doi.org/10.1175/2009WAF2222289.1

[14] Bendix, J., 2002. A satellite-based climatology of fog and low-level stratus in Germany and adjacent areas. Atmospheric Research. 64, 3–18. DOI: https://doi.org/10.1016/S0169-8095(02)00075-3

[15] Ahn, M., Sohn, E., Hwang, B., 2003. A new algorithm for sea fog/stratus detection using GMS-5 IR data. Advances in Atmospheric Sciences. 20, 899–913. DOI: https://doi.org/10.1007/BF02915513

[16] Bendix, J., Cermak, J., Thies, B., 2004. New Perspectives in Remote Sensing of Fog and Low Stratus—TERRA/AQUA-MODIS and MSG. In Proceedings of the 3rd International Conference on Fog, Fog Collection and Dew, Cape Town, South Africa, 11–15 October 2004; pp. G2.1–G2.4.

[17] Bendix, J., Thies, B., Cermak, J., et al., 2005. Ground fog detection from space based on MODIS daytime data—A feasibility study. Weather and Forecasting. 20, 989–1005. DOI: https://doi.org/10.1175/WAF886.1

[18] Ellrod, G.P., Lindstrom, S., 2006. Performance of satellite fog detection techniques associated with major fog-related highway accidents. Environmental Science and Engineering. Available from: https://www.semanticscholar.org/paper/Performance-of-Satellite-Fog-Detection-Techniques-Ellrod-Lindstrom/af145aee2c5f94a24b56ecb0f42c3fb8734193d6

[19] Yoo, J., Jeong, M., Yoo, H., et al., 2006. Fog sensing over the Korean Peninsula derived from satellite observation of MODIS and GEOS-9. Korean Journal of Remote Sensing. 22, 373–377.

[20] Choudhury, S., Rajpal, H., Saraf, A.K., et al., 2007. Technical Note: Mapping and forecasting of North Indian winter fog: An application of spatial technologies. International Journal of Remote Sensing. 28, 3649–3663.

[21] Ward, B., Croft, P.J., 2008. Use of GIS to examine winter fog occurrences. Electronic Journal of Operational Meteorology. 33. Available from: https://www.semanticscholar.org/paper/Use-of-GIS-to-Examine-Winter-Fog-Occurrences-Ward-Croft/6ce9e6015978c5b31aa77c56b90b6b1f843ec60f

[22] Zhang, J., Lu, H., Xia, Y., et al., 2018. Deep Convolutional neural network for fog detection. In Intelligent Computing Theories and Application. Springer: Cham, Switzerland. DOI: https://doi.org/10.1007/978-3-319-95933-7_1

[23] Kopecká, J., Kopecký, D., Štursa, D., et al., 2026. Estimation of atmospheric visibility by deep learning model using multimodal dataset. Knowledge-Based Systems. 331, 114732. DOI: https://doi.org/10.1016/j.knosys.2025.114732

[24] Castillo-Botón, C., Casillas-Pérez, D., Casanova-Mateo, C., et al., 2022. Machine learning regression and classification methods for fog events prediction. Atmospheric Research. 272, 106157. DOI: https://doi.org/10.1016/j.atmosres.2022.106157

[25] Shi, X., Chen, Z., Wang, H., et al., 2015. Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. arXiv preprint. arXiv:1506.04214.

[26] Wagenbrenner, N.S., Forthofer, J.M., Page, W.G., et al., 2019. Development and Evaluation of a Reynolds-Averaged Navier–Stokes Solver in WindNinja for Operational Wildland Fire Applications. Atmosphere. 10, 672. DOI: https://doi.org/10.3390/atmos10110672

[27] Xiang, Y., Zhang, Q., Wang, M., et al., 2025. A hybrid forecasting approach for sea fog based on machine learning. Artificial Intelligence for the Earth Systems. 5, e250013. DOI: https://doi.org/10.1175/AIES-D-25-0013.1

[28] Pagowski, M., Gultepe, I., King, P., 2004. Analysis and modeling of an extremely dense fog event in southern Ontario. Journal of Applied Meteorology and Climatology. 43, 3–16. DOI: https://doi.org/10.1175/1520-0450(2004)043<0003:AAMOAE>2.0.CO;2

[29] Hinton, G.E., Salakhutdinov, R.R., 2006. Reducing the dimensionality of data with neural networks. Science. 313, 504–507. DOI: https://doi.org/10.1126/science.1127647

[30] Palvanov, A., Cho, Y.I., 2019. VisNet: Deep convolutional neural networks for forecasting atmospheric visibility. Sensors. 19, 1343. DOI: https://doi.org/10.3390/s19061343

[31] Vaswani, A., Shazeer, N., Parmar, N., et al., 2017. Attention is All You Need. arXiv preprint. arXiv:1706.03762.

[32] Steeneveld, G.J., Ronda, R.J., Holtslag, A.A.M., 2015. The challenge of forecasting the onset and development of radiation fog using mesoscale atmospheric models. Boundary-Layer Meteorology. 154, 265–289. DOI: https://doi.org/10.1007/s10546-014-9973-8

[33] Ranjan, P., Akshay, V., Kumar, P.R.G., et al., 2025. Improving weather predictions accuracy with a hybrid LSTM-Transformer model: A study on Indian climate. IEEE Access. 13, 169014–169025. DOI: https://doi.org/10.1109/ACCESS.2025.3614473

[34] Miao, K.C., Han, T.T., Yao, Y.Q., et al., 2020. Application of LSTM for short term fog forecasting based on meteorological elements. Neurocomputing. 408, 285–291.

[35] Lu, Z., Zheng, C., Yang, T., 2020. Application of offshore visibility forecast based on temporal convolutional network and transfer learning. Computational Intelligence and Neuroscience. 8882279. DOI: https://doi.org/10.1155/2020/8882279

[36] Kamangir, H., Collins, W., Tissot, P., et al., 2021. FogNet: A multiscale 3D CNN with double-branch dense block and attention mechanism for fog prediction. Machine Learning with Applications. 5, 100038. DOI: https://doi.org/10.1016/j.mlwa.2021.100038

[37] Chen, K., Zhou, Y., Ren, T., et al., 2024. Short-term sea fog area forecast: A new data set and deep-learning pixel-level model. Journal of Geophysical Research. DOI: https://doi.org/10.1029/2024JH000230

[38] Qu, Y., Fang, Y., Ji, S., et al., 2024. Deep learning-based atmospheric visibility detection. Atmosphere. 15,

1394. DOI: https://doi.org/10.3390/atmos15111394
[39] Peláez-Rodríguez, C., Pérez-Aracil, J., Casanova-Mateo, C., et al., 2023. Efficient prediction of fog-related low-visibility events with Machine Learning and evolutionary algorithms. Atmospheric Research. 295, 106991. DOI: https://doi.org/10.1016/j.atmosres.2023.106991
[40] Salcedo-Sanz, S., Guijo-Rubio, D., Pérez-Aracil, J., et al., 2025. Artificial intelligence-based methods and algorithms in fog and atmospheric low-visibility forecasting. Atmosphere. 16, 1073. DOI: https://doi.org/10.3390/atmos16091073
[41] Smith, D.K.E., Reka, S., Dorling, S.R., et al., 2024. Forecasts of fog events in northern India dramatically improve when weather prediction models include irrigation effects. Communications Earth & Environment. 5, 141. DOI: https://doi.org/10.1038/s43247-024-01314-w
[42] Shankar, A., Kumar, A., Sinha, V., 2024. Machine learning approach in the prediction of fog: An early warning system. MAUSAM. 75, 1039–1050. DOI: https://doi.org/10.54302/mausam.v75i4.5919
[43] Ortega, L.C., Otero, L.D., Solomon, M., et al., 2023. Deep learning models for visibility forecasting using climatological data. International Journal of Forecasting. 39, 992–1004.
[44] Singh, A., Maheskumar, R.S., Iyengar, G.R., 2022. A diagnostic method for fog forecasting using numerical weather prediction (NWP) model outputs. Journal of Atmospheric Science Research. 5(4), 10–19. DOI: https://doi.org/10.30564/jasr.v5i4.5068
[45] Bartok, J., Bott, A., Gera, M., 2012. Fog prediction for road traffic safety in a coastal desert region. Boundary-Layer Meteorology. 145, 485–506. DOI: https://doi.org/10.1007/s10546-012-9750-5
[46] Safai, P.D., Ghude, S., Pithani, P., et al., 2019. Two-way relationship between aerosols and fog: A case study at IGI Airport, New Delhi. Aerosol and Air Quality Research. 19(1), 71–79. DOI: https://doi.org/10.4209/aaqr.2017.11.0542
[47] Arun, S.H., Singh, C., John, S., et al., 2022. A study to improve the fog/visibility forecast at IGI Airport, New Delhi during the winter season 2020–2021. Journal of Earth System Science. 131, 124. DOI: https://doi.org/10.1007/s12040-022-01874-5
[48] Goswami, S., Chaudhuri, S., Das, D., et al., 2020. Adaptive neuro-fuzzy inference system to estimate the predictability of visibility during fog over Delhi, India. Meteorological Applications. 27, e1900. DOI: https://doi.org/10.1002/met.1900
[49] Gal, Y., Ghahramani, Z., 2015. Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. arXiv preprint. arXiv:1506.02142.
[50] MacKay, D.J., 1992. A practical Bayesian framework for backpropagation networks. Neural Computation. 4, 448–472.
[51] Neal, R.M., 2012. Bayesian Learning for Neural Networks. Springer: New York, NY, USA.
[52] Blundell, C., Julien, C., Koray, K., et al., 2015. Weight uncertainty in neural network. arXiv preprint. arXiv:1505.05424.
[53] Rasmussen, C.E., Williams, C.K.I., 2006. Gaussian Processes for Machine Learning. MIT Press: Cambridge, MA, USA.
[54] Gneiting, T., Katzfuss, M., 2014. Probabilistic forecasting. Annual Review of Statistics and Its Application. 1, 125–151. DOI: https://doi.org/10.1146/annurev-statistics-062713-085831
[55] Raftery, A.E., Gneiting, T., Balabdaoui, F., et al., 2005. Using Bayesian model averaging to calibrate forecast ensembles. Monthly Weather Review. 133, 1155–1174. DOI: https://doi.org/10.1175/MWR2906.1
[56] Bishop, C.M., Nasrabadi, N.M., 2006. Pattern Recognition and Machine Learning. Springer: New York, NY, USA.
[57] Pinson, P., 2013. Wind energy: Forecasting challenges for its operational management. Statistical Science. 28, 564–585. DOI: https://doi.org/10.1214/13-STS445
[58] Boukabara, S.A., Krasnopolsky, V., Penny, S.G., et al., 2021. Outlook for exploiting artificial intelligence in the earth and environmental sciences. Bulletin of the American Meteorological Society. 102, E1016–E1032. DOI: https://doi.org/10.1175/BAMS-D-20-0031.1
[59] Ait Ouadil, K., Idbraim, S., Bouhsine, T., et al., 2024. Atmospheric visibility estimation: A review of deep learning approach. Multimedia Tools and Applications. 83, 36261–36286. DOI: https://doi.org/10.1007/s11042-023-16855-z
[60] Kuleshov, V., Fenner, N., Ermon, S., 2018. Accurate uncertainties for deep learning using calibrated regression. arXiv preprint. arXiv:1807.00263.
[61] Rahaman, R., Theiry, A.H., 2020. Uncertainty quantification and deep ensembles. arXiv preprint. arXiv:2007.08792. DOI: https://doi.org/10.48550/arXiv.2007.08792
[62] Bajaj, K., Mannam, U., Deshpande, P., et al., 2024. Forecasting of fog index and prediction interval using Bayesian methods. In Proceedings of the 8th International Conference on Data Science and Management of Data (12th ACM IKDD CODS and 30th COMAD), Jodhpur, India, 18–21 December 2024; pp. 270–278. DOI: https://doi.org/10.1145/3703323.3703738