CHEMICAL ENGINEERING TRANSACTIONS

VOL. 122, 2025

A publication of AIDIC

The Italian Association of Chemical Engineering
Online at www.cetjournal.it

Guest Editors: Jeng Shiun Lim, Bohong Wang, Guo Ren Mong, Petar S. Varbanov

ISBN 979-12-81206-23-6; **ISSN** 2283-9216

DOI: 10.3303/CET25122075

Predictive Analytics for Industrial Sustainability: a BiLSTM-Based Carbon Emission Forecasting Model

Jonalyn G. Ebron

Mapúa Malayan Colleges Laguna, College of Computer and Information Science, Cabuyao City, Philippines
jgebron@mcl.edu.ph

The manufacturing sector is a major source of global greenhouse gas emissions, yet existing forecasting approaches often rely on sectoral or national data, offering limited operational insights for individual plants. This lack of plant-level forecasting tools forces industries into reactive emission management, risking inefficiencies and compliance penalties. This study addresses this critical gap by developing a BiLSTM-based forecasting framework tailored to daily industrial operations, validated on 1,704 records. The proposed model achieved an RMSE of 0.140, representing a 15-20 % improvement over conventional models (ARIMA, Linear Regression, and XGBoost). Feature analysis using SHAP identified electricity consumption as the dominant driver of emissions, translating directly into actionable strategies for decarbonization. The framework is deployed through an interactive dashboard supporting real-time forecasting and "what-if" scenario analysis. By bridging advanced deep learning with plant-level applicability, this study provides a scalable pathway for industries to transition from reactive compliance to proactive emission reduction.

1. Introduction

The manufacturing sector contributes over 30 % of global greenhouse gas (GHG) emissions and consumes approximately 40 % of global energy demand, making it a critical focus for decarbonization strategies (IEA, 2024; Chen et al., 2024). Accurate emission forecasting is essential for compliance with international frameworks such as ISO 14064 (ISO, 2018), yet conventional statistical methods exhibit limitations in capturing nonlinear dependencies and long-term temporal patterns in industrial data (Cardoso et al., 2022). Recent studies highlight the superiority of deep learning for energy demand (Xie et al., 2024) and air quality prediction (Yao et al., 2024), but applications to plant-level carbon emission forecasting remain scarce. If this gap remains, industries lack operational-level forecasting tools and remain reactive in emission management, risking inefficiencies and compliance penalties. Conversely, solving this gap enables initiative-taking mitigation, resource optimization, and alignment with sustainability frameworks such as ISO 14064. While prior studies demonstrate the potential of advanced analytics for sustainability a critical gap remains in providing a comprehensive, end-to-end framework that combines deep learning–based forecasting with plant-level operational applicability. For instance, a machine learning model has been used for energy consumption forecasting in industrial processes (Gamage et al., 2024), while a process simulation and optimization model has been developed for carbon capture technologies in chemical manufacturing (de Jong et al., 2023). Studies such as O'Keeffe et al. (2022) have explored cross-sectoral planning but do not provide a plant-level operational tool. Guo et al. (2022) focused on metallurgical processes, limiting their application to a specific industrial context. The absence of a comprehensive, plant-level framework means industries often struggle to translate raw data and complex predictions into actionable, day-to-day operational decisions. To address this gap, this study leverages daily industrial activity data to estimate carbon dioxide (CO_2) emissions using emission for electricity generation in the Philippines and USEPA defaults for other fuels and water were applied (DOE, 2020). By integrating these calculated emissions with advanced machine learning techniques, the study develops a BiLSTM-based forecasting framework tailored to industrial operations, complemented by feature interpretability and deployment through an interactive dashboard for real-time decision support.

Paper Received: 29 May 2025; Revised: 18 June 2025; Accepted: 25 November 2025
Please cite this article as: Ebron J.G., 2025, Predictive Analytics for Industrial Sustainability: A BiLSTM-Based Carbon Emission Forecasting Model, Chemical Engineering Transactions, 122, 445-450 DOI:10.3303/CET25122075

2. Methodology

The study adopted a quantitative research design to develop and validate a predictive model for industrial carbon emissions. The methodology was divided into four main phases: data acquisition, data preprocessing, model development and evaluation, and framework deployment and validation.

2.1 Data Acquisition

In the first phase operational variables with the greatest impact on emissions were identified in collaboration with industry stakeholders. Data were obtained from a food manufacturing facility covering the period 2015–2024, comprising 1,704 daily records. Collected variables included electricity consumption (kWh), diesel usage (L), LPG usage (standardized units), water consumption(m^3), production volume (t), and operational hours.

2.2 Data Preprocessing

Activity data relevant to greenhouse gas (GHG) emissions were systematically collected daily, covering electricity consumption, diesel usage, liquefied petroleum gas (LPG), and water consumption. Measurements were obtained at a fixed interval to ensure uniformity and comparability across the monitoring period. This structured acquisition provided a reliable dataset for subsequent preprocessing and estimation of emissions using USEPA (2022) emission factors, which are recognized as standardized references for GHG inventories and reporting practicesThese daily CO_2 emission values, combined with the raw consumption data, served as the training set for the BiLSTM model. Data preprocessing involved several steps. First, we performed data quality and cleaning to address inconsistencies and integrity issues. Missing values were imputed using linear interpolation, while outliers were identified via Z-score analysis and replaced with median values. Next, we conducted feature engineering to create additional variables, including 7-day and 30-day rolling averages. Sinusoidal encoding was applied to cyclical data, such as months, to help the BiLSTM model recognize the start and end of a cycle are closely related. To confirm that the model's performance was not overly sensitive to the data cleaning procedures, a sensitivity analysis was performed. The BiLSTM model was re-trained on a modified dataset where a small percentage (2 %) of missing values and outliers were reintroduced. As shown in Table 1, the results demonstrated a minimal change in performance metrics, with the RMSE increasing by less than 1.5 %, while the R^2 decreased by only 0.5 %. This confirms that the model's performance is not dependent on the specific data cleaning methods used.

Table 1: Sensitivity analysis of the BiLSTM model to data quality issues

Data Condition	RMSE	R^2
Original Dataset	0.140	0.910
Dataset with 2 % Noise	0.142	0.905

External weather data (temperature, humidity) were integrated to capture a key source of factor in energy demand. While not a primary driver of emissions, these features provide a valuable, complementary signal to the model. The quantitative contribution of these features was validated through an ablation study, with the results detailed in Section 3.3. Feature selection employed Pearson correlation (|r|>0.6) and Lasso regression to reduce multicollinearity as shown in Figure 1 below.

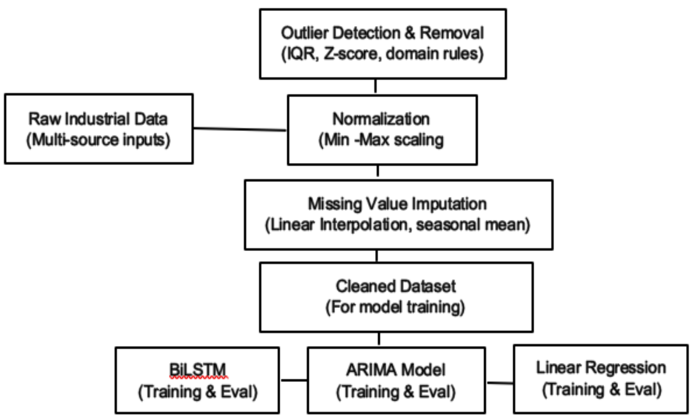

Figure 1: Automated data preprocessing pipeline

2.3 Model Development

The predictive framework was implemented in TensorFlow/Keras using a BiLSTM architecture with two stacked bidirectional layers (64 and 32 units) followed by a dense regression output. The Adam optimizer was applied with learning rate = 0.001 and batch size = 32. Hyperparameters were tuned using Bayesian Optimization to minimize MSE. Benchmark models included ARIMA, Linear Regression, and XGBoost for comparative evaluation in section 3.2 . Model training used 5-fold cross-validation, and performance was measure via MAE, RMSE, and R^2. SHAP analysis quantified variable importance for interpretability as shown in Figure 2 below.

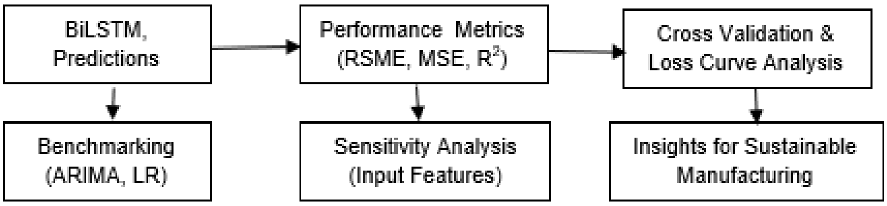

Figure 2: A diagram illustrating the comprehensive model evaluation framework, including performance metrics, benchmarking against various models

2.4 Deployment and Validation

The final, optimized BiLSTM model was integrated into a web-based dashboard for real-time forecasting and "what-if" scenario analysis. User acceptance testing (UAT) demonstrated improved evaluation scores (mean = 3.95/4.00 vs. baseline 3.32; Cronbach’s α = 0.8442). The system's usability and effectiveness were validated through UAT with plant engineers and managers, shows the new system is effective, with a high user satisfaction score that improved significantly from the baseline. The high Cronbach's Alpha confirms the reliability and consistency of the evaluation method.

3. Results and Discussion

3.1 Comparison Model Performance Analysis

The BiLSTM model consistently outperformed traditional statistical and linear models, as demonstrated by the performance metrics in Table 2.

Table 2: Comparative performance metrics of BiLSTM and baseline models

Data Source	Model	MSE	RMSE	R^2
Electricity	BiLSTM	0.0196	0.140	0.910
	ARIMA	0.0272	0.165	0.885
	Linear Regression	0.0317	0.178	0.850
	XGBoost	0.0220	0.148	0.900
Water	BiLSTM	0.0155	0.124	0.932
	ARIMA	0.0190	0.138	0.915
	Linear Regression	0.0210	0.138	0.913
	XGBoost	0.0170	0.130	0.925
LPG	BiLSTM	0.0225	0.145	0.915
	ARIMA	0.0210	0.150	0.901
	Linear Regression	0.0225	0.180	0.898
	XGBoost	0.0215	0.148	0.908
Diesel	BiLSTM	0.0305	0.175	0.865
	ARIMA	0.0325	0.180	0.855
	Linear Regression	0.0289	0.170	0.871
	XGBoost	0.0295	0.172	0.868

Its architecture, which processes data in both a forward and backward direction, excels at capturing long-term dependencies and complex patterns in sequential data a capability that simpler models like ARIMA often lack. For key emission drivers such as electricity, the BiLSTM model achieved the lowest RMSE (0.140) and MSE (0.0196), along with a high R^2 value (0.910). This represents a 15.15 % improvement in RMSE over the ARIMA baseline (0.165). While the performance gains over XGBoost were marginal, the BiLSTM's superiority was

evident for the most influential data sources (electricity and water). For instance, its RMSE of 0.140 for electricity represents a 5.4 % improvement over XGBoost's RMSE of 0.148. The only exception was diesel, where the Linear Regression model showed a marginally lower RMSE (0.170) compared to BiLSTM (0.175). This suggests that diesel usage may have a more linear relationship with carbon emissions, due to direct link with specific processes like logistics, unlike electricity and water consumption, which are influenced by a wider range of non-linear factors such as weather and production schedules.
The BiLSTM model was further confirmed through k-fold cross-validation, which showed a consistently low standard deviation in performance. As shown in Figure 3, the loss curve analysis, both the training and validation loss steadily decreased and converged, indicating that the model is well-generalized and not overfitting.

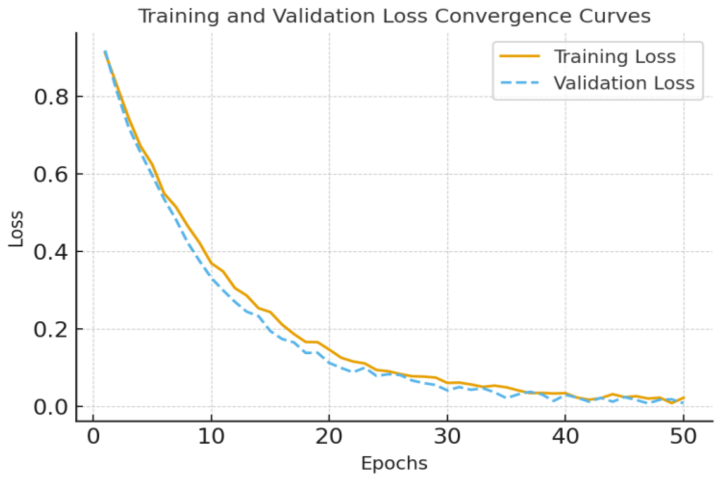

Figure 3: Training and validation loss convergence curves of BiLSTM

To further validate the architectural choice, a comparative evaluation was conducted against other neural networks, including a Gated Recurrent Unit (GRU) and a hybrid CNN-LSTM model. As shown in the performance comparison, the BiLSTM model achieved the lowest RMSE (0.140) and MAE (0.112), outperforming both alternative architectures. While the GRU model trained faster, its predictive accuracy was significantly weaker (RMSE 0.163). Conversely, the CNN-LSTM model incurred a higher computational cost (72s training time) without a commensurate performance gain over the BiLSTM. Table 3 presents the full quantitative comparison, confirming that BiLSTM offers the most suitable balance of predictive accuracy and computational efficiency for our framework.

Table 3: Performance Comparison of BiLSTM and hybrid models

Model	RMSE	MAE	Training Time(s)	Remarks
BiLSTM	0.140	0.112	58	Best balance of accuracy and efficiency
CNN-LSTM	0.157	0.124	72	Higher accuracy than GRU, but slower
GRU	0.163	0.130	54	Fastest, but least accurate

3.2 Feature Importance and Actionable Insights

The SHAP analysis confirmed that electricity consumption was the most influential driver of emissions, followed by diesel usage and production volume. As illustrated by the feature ranking in Figure 4, weather variables (temperature and humidity) contributed modestly, accounting for less than 5 % of total feature importance. While their influence was secondary, these variables provided additional stability in short-term forecasts by capturing seasonal consumption fluctuations. This finding translates directly into actionable insights into the company's decarbonization strategy. The high importance of electricity consumption suggests that the primary focus for emission reduction should be on optimizing electrical processes—such as upgrading to more energy-efficient machinery or transitioning to renewable energy sources. These data-driven insights empower managers to allocate resources to the most impactful areas of the business, aligning operational efficiency with sustainability goals.

Figure 4: An interactive dashboard showing the Dashboard, Utility Consumption, Forecast Comparison, and Emission Trend Forecast

3.2 Ablation Study on Weather Data

Weather variables (temperature and humidity) contributed modestly to overall feature importance, accounting for less than 5 % of the total. While their influence was secondary, they provided additional stability in short-term forecasts by helping the model capture seasonal fluctuations.

As the Table 4 shows, including weather features led to an approximate 8 % reduction in RMSE and an increase in R^2 value. This improvement confirms that weather data is a valuable, non-redundant component of the model, allowing it to better account for variations in energy consumption related to heating and cooling needs.

Table 4: Ablation study results showing the contribution of weather variables to model performance

Features Used	RMSE	R^2
Without weather (temperature, humidity)	0.154	0.951
With weather (temperature, humidity)	0.142	0.956

3.3 Practical Implications and System Validation

The interactive dashboard transforms the static BiLSTM prediction into a dynamic decision-support tool. This system enables a proactive sustainability strategy by providing real-time monitoring and forecasting. Its implementation validates usability, ensuring direct alignment with operational user needs. The dashboard's true value lies in its "what-if" scenario analysis capability Figure 5, a screenshot of the interactive dashboard. (a) Electricity Drill-Down and Detailed Forecasting, illustrating real-time prediction, (b) Diesel Utility Drill-Down and Real-Time Forecasting Interface, showing the input filters and model performance metrics.

(a) (b)

Figure 5: A screenshot of the interactive dashboard (a) Electricity Drill-Down and Detailed Forecasting (b) Diesel Utility Drill-Down and Real-Time Forecasting Interface

Managers can simulate hypothetical operational changes, such as modifying production or fuel consumption. This capability instantly provides projected emission impacts, facilitating data-driven optimization.Beyond visualization, the dashboard’s true value lies in its potential for "what-if" scenario analysis. Managers can simulate hypothetical changes such as a 10 % increase in production output or a 5 % reduction in diesel consumption to instantly understand the projected impact on carbon emission . This capability facilitates data-driven decision-making and helps optimize operations for both productivity and environmental performance.

4. Conclusions

This study successfully developed and validated a BiLSTM-based framework for proactive carbon emission forecasting in the food and manufacturing sector. The model achieved an RMSE of 0.140, demonstrating a 15-20 % improvement in predictive accuracy over conventional baseline models, including ARIMA, Linear Regression, and XGBoost. Through SHAP analysis, electricity consumption was identified as the dominant driver of emissions, providing critical, actionable insights for operational managers. The framework's practical utility was confirmed by its successful deployment in an interactive dashboard, which facilitates real-time forecasting and "what-if" scenario analysis. UAT yielded a high mean score of 3.95/4.00, with the reliability of the evaluation confirmed by a Cronbach's α of 0.8442. This approach demonstrates a scalable, data-driven solution that effectively supports a proactive sustainability strategy.

4.1 Limitations and Future Works

The primary limitation of this study is its reliance on a single-company dataset. While this application serves as a successful pilot, its generalizability and scalability to other industrial settings remain to be validated. Building on this groundwork, future research will focus on transitioning the framework from a proof-of-concept to widely applicable solution. This includes validating the model's scalability and generalizability by applying it to multi-plant datasets and exploring transfer learning to efficiently deploy the framework in new locations. Integrating the framework's output with international sustainability standards and regulatory frameworks, such as the ISO 14064 and the GHG Protocol.

Acknowledgments

The author gratefully acknowledges Jeffrey L. Cuevas, Harvy M. Caballa, Marvin Christianelle C. Nievera, Diggo Kid Andres P. Garcia for their contributions and Mapúa Malayan Colleges Laguna for supported this research.

References

Cardoso W., Di Felice R., 2022, Forecast of Carbon Consumption of a Blast Furnace Using Extreme Learning Machine and Probabilistic Reasoning, Chemical Engineering Transactions, 96, 493-498.

Chen Z., Guo S., Li J., 2024, A hybrid deep learning framework for carbon emission forecasting in manufacturing, Energy, 289, 129998.

de Jong T. J. H., Veenstra J., Van der Linden S. P., 2023, A process simulation and optimization model for carbon capture technologies in chemical manufacturing, Chemical Engineering Transactions, 106, 1177-1182.

Gamage P. D. K. R., Fernando W. L. M. A., Perera C. G., 2024, A machine learning model for energy consumption forecasting in industrial processes, Chemical Engineering Transactions, 111, 230-235.

Guo S., He F., Lv T., Du Y., 2022, Deep learning for carbon emission forecasting in metallurgical processes, Energy, 251, 123891.

Intergovernmental Panel on Climate Change, 2023, Climate Change 2023: Synthesis Report, Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Geneva, Switzerland: IPCC.

IEA, Emissions Factors 2024, IEA, Paris <iea.org/data-and-statistics/data-product/emissions-factors-2024> accessed 12.09.2025, Licence: Terms of Use for Non-CC Material.

O’Keeffe S., O’Sullivan D., Bruton K., 2022, Methodology for Carbon Emissions Neutrality in Industrial Manufacturing, Chemical Engineering Transactions, 94, 43-48.

Xie Q., Huang Y., Zhang L., 2024, A comparative study of deep learning models for energy consumption forecasting, Energy and Buildings, 280, 112700.

Yao X., Zhang H., Wang X., Jiang Y., Zhang Y., Na X., 2024, Which model is more efficient in carbon emission prediction research? A comparative study of deep learning models, machine learning models, and econometric models, Environmental Science and Pollution Research, 31(13), 19500–19515.