AIMS Mathematics, 11(4): 9655–9685.
DOI: 10.3934/math.2026400
Received: 19 January 2026
Revised: 06 March 2026
Accepted: 19 March 2026
Published: 10 April 2026

https://www.aimspress.com/journal/Math

Research article

A dual encounter logarithmic path neural network for precision carbon emission monitoring and mitigation

Mohammad Barr[1], Tawfeeq Shawly[2], Ahmed A. Alsheikhy[3,*] and Sahbi Boubaker[3]

[1] Department of Electrical Engineering, College of Engineering, Northern Border University, Arar 91431, Saudi Arabia; Email: mohammed.barr@nbu.edu.sa
[2] Department of Electrical Engineering, Faculty of Engineering at Rabigh, King Abdulaziz University, Jeddah 21589, Saudi Arabia; Email: tshawly@kau.edu.sa
[3] Department of Computer and Network Engineering, College of Computer Science and Engineering, University of Jeddah, Jeddah 21959, Saudi Arabia; Email: aaalsheikhy@uj.edu.sa, sboubaker@uj.edu.sa

* **Correspondence:** Email: aaalsheikhy@uj.edu.sa.

Abstract: Monitoring and reducing carbon footprints are crucial for achieving sustainability and effectively tackling climate change. Given the urgent global need to address climate change and to participate in Saudi Arabia's Vision 2030, we present DeepCarbonNet (EcoNet), a novel deep learning (DL) framework to monitor, analyze, and reduce carbon emissions. The framework is built using a new Dual Encounter Logarithmic Path Neural Network (DELPNN) architecture, including a novel Spatial Encounter Pathway (SEP), which processes high-resolution satellite images through a Logarithmic Convolutional Encoder (LCE) to extract multi-scale spatial features, and a new Temporal Encounter Pathway (TEP), which processes sequential Internet of Things (IoT) sensor and energy consumption data via a Gated Logarithmic Recurrent Unit (GLRU) and a central Feature Fusion Operator (FFO) that integrates the spatial and temporal features using cross-attention mechanisms and projects them into a logarithmic latent space to capture intricate, non-linear emission dynamics. This approach enables the precise capture of spatial and temporal dependencies within carbon emission data, achieving an outstanding level of accuracy. The simulation experimental results demonstrated that EcoNet attains a high accuracy of 98.7% in estimating carbon footprints after training the model on two public datasets. Furthermore, the model employed a reinforcement learning (RL)-based optimization strategy, enabling a 29.4% reduction in emissions through adaptive mitigation techniques. EcoNet was designed to adapt to changing conditions and promote environmental sustainability

continuously. Beyond monitoring, EcoNet achieved a 32.8% improvement in energy efficiency. Additionally, the framework demonstrated robust performance across weather conditions, with 97.0-98.7% accuracy and an accuracy of emission intensities between 94.2–99.1%. These results showed that EcoNet is a solution for artificial intelligence (AI)-driven environmental sustainability, which offers immediate practical value for industrial monitoring, smart city management, logistic services to reduce fuel consumption, and national sustainability programs.

Keywords: DeepCarbonNet; climate change; sustainability; carbon emission; carbon footprint monitoring; emission reduction optimization; real-time monitoring; logistics
Mathematics Subject Classification: 68T07, 68T10, 68T20

1. Introduction

As global economies grow and industrial activities increase, the need for tracking and reducing carbon footprints has reached a critical point [1]. Carbon emissions, which come from various sources, such as fossil fuel burning, industrial processes, transportation, and deforestation, play a major role in climate change, resulting in increased global temperatures, severe weather events, and environmental deterioration [1]. The United Nations Intergovernmental Panel on Climate Change (IPCC) has indicated that a 45% reduction in global carbon emissions is essential by 2030 to alleviate the most severe consequences of climate change [1]. Nevertheless, conventional methods of carbon tracking often depend on manual data collection, outdated emissions models, and sporadic reporting, rendering them ineffective in addressing contemporary environmental issues [1,2]. Developments in artificial intelligence (AI), deep learning (DL), and edge computing have paved the way for automated, scalable, and highly precise solutions for real-time assessment and reduction of carbon footprints [3–6]. AI-driven carbon tracking facilitates continuous monitoring, swift anomaly detection, and predictive emissions modeling, empowering industries, governments, and policymakers to make informed, data-driven choices. The synergy of AI with smart city initiatives, industrial refineries, and renewable energy systems further refines carbon management strategies, enhancing resource efficiency while minimizing environmental impact. The global shift toward carbon-neutral economies has prompted a significant increase in technological innovations focused on tracking and mitigating carbon footprints [1]. A variety of AI-driven approaches have emerged, utilizing IoT-based emissions monitoring, machine learning (ML)-enhanced predictive analytics, and blockchain-enabled carbon credit systems [3–6].

Saudi Arabia's Vision 2030 prioritizes the integration of green technologies and AI to meet its carbon reduction objectives [2]. The Saudi Green Initiative aims to decrease carbon emissions by 278M tons annually, positioning the nation as a frontrunner in sustainable energy and environmental stewardship [7,8]. The implementation of AI for tracking carbon footprints is essential for adhering to global sustainability standards, including the Paris Agreement and the United Nations Sustainable Development Goals (SDGs) [9–12].

The evolution of DL methodologies has transformed the landscape of carbon footprint monitoring by facilitating precise, automated, and scalable emissions assessments [11–14]. A variety of AI frameworks have been developed to tackle the intricacies of carbon tracking, integrating computer vision, natural language processing (NLP), and reinforcement learning (RL) to improve monitoring, forecasting, and optimization efforts [14–19]. Among the most proficient DL models employed for

carbon footprint evaluation are Vision Transformers (ViTs), Graph Neural Networks (GNNs), Recurrent Neural Networks (RNNs), and models based on RL [18]. These frameworks deliver accurate and dynamic insights into carbon emissions, enabling real-time monitoring, prediction, and informed decision-making. ViTs are especially advantageous for interpreting satellite imagery and aerial surveillance data, facilitating the identification of carbon emissions from industrial areas, deforested regions, and transportation systems. ViTs surpass conventional Convolutional Neural Networks (CNNs) by utilizing self-attention mechanisms, which enable them to analyze extensive environmental datasets with greater precision. Likewise, GNNs play a crucial role in modeling the intricate relationships among emission sources, energy infrastructures, and environmental variables. GNNs can effectively delineate the interconnected elements that contribute to carbon emissions, offering valuable insights into supply chain emissions, industrial waste management, and urban carbon footprints. RNNs and Long Short-Term Memory (LSTM) networks are widely used to analyze time-series data for predicting carbon footprints.

1.1. The problem statement

Despite the availability of monitoring technologies, accurately quantifying and mitigating carbon emissions remains a significant challenge for nations pursuing sustainability targets. Current methodologies mostly rely on periodic manual reporting, diverse sensor networks, and an inability to capture real-time emission fluctuation data, which results in analysis failure for the detection of anomalies, inefficient resource allocation for mitigation efforts, and an inability to respond to changing environmental conditions. Moreover, the government of Saudi Arabia has launched its Vision 2030 to use renewable energy and reduce fossil fuel by 50% by 2030 in its massive projects, such as Neom and the Red Sea Project. These projects require intelligent environmental management systems to maintain low-carbon urban operations. However, the approaches lack the capabilities and response mechanisms for such complex and real-time applications.

We bridge this gap by developing a system that can simultaneously process heterogeneous environmental data streams, identify emission patterns, and generate actionable insights for immediate mitigation interventions. To participate in Vision 2030, we introduce EcoNet, a novel DL framework that is designed for real-time carbon monitoring and adaptive emission reduction.

1.2. Objectives and motivations

Our major motivations for this study are:

1- Climate Change Mitigation: The implementation of AI-based carbon tracking is essential for meeting the global objective of achieving net-zero emissions by 2050.
2- Industry 4.0 and Smart Manufacturing: AI models operating in real-time enable industries to continuously monitor and effectively reduce carbon emissions.
3- Economic Incentives: The utilization of AI in carbon tracking can enhance carbon credit trading, projected to reach a market value of $250B by 2030.

Our objectives of this research are summarized as follows:

I. Develop a DL model architecture for the real-time evaluation of carbon footprints.
II. Incorporate diverse data sources, including IoT sensors, satellite imagery, and industrial reports, to improve tracking precision.

III. Attain a minimum accuracy rate of 98.7% in forecasting carbon emissions.

IV. Utilize AI-enhanced RL to refine carbon reduction strategies, aiming for a 35% reduction in emissions.

V. Ensure that AI-led sustainability initiatives are in line with Vision 2030 objectives for industrial transformation and environmental stewardship.

VI. Supports the Saudi Green Initiative, which aims for a reduction of 278M tons of CO_2 annually.

VII. Facilitates automated compliance monitoring for industries transitioning to net-zero emissions.

VIII. Improves AI-driven energy efficiency initiatives, minimizing energy waste in smart cities and infrastructure developments.

Therefore, our implemented scheme contributes to environmental science by addressing the challenges of real-time and multi-source data integration for monitoring of carbon emission and mitigation.

1.3. Contributions

This study offers multiple contributions, which are:

1- Development of the Carbon Monitoring Framework: Presents a novel DL architecture (DELPNN) that combines SEP, TEP, GLRU, and FFO to effectively monitor and evaluate carbon emissions across sectors.

2- Multimodal Data Fusion for Enhanced Accuracy: Incorporates four distinct data sources, satellite imagery, IoT sensor data, industrial emission records, and climate reports, to enhance the estimation of carbon footprints, achieving an accuracy rate of 98.7%, which surpasses traditional models by 14.2%.

3- AI-Powered RL for Carbon Reduction: Adopts an adaptive RL-based optimization approach, resulting in a 29.4% decrease in emissions by continuously modifying mitigation strategies based on real-time environmental data.

4- Adaptive Solution for Smart Cities and Industrial Sustainability: Engineered for smooth integration with smart city frameworks, industrial sustainability programs, and AI-driven energy management systems, aligning with Saudi Arabia's Vision 2030 sustainability objectives.

The organization of this paper is outlined as follows: In Subsection 1.4, we provide a review of the literature, while in Section 2, we detail the proposed methodology. In Section 3, the experimental settings and results are discussed, followed by an analysis of the implications and limitations in Section 4. Finally, in Section 5, we offer conclusions along with recommendations for future research.

1.4. Literature review

A range of techniques, including ML, DL, RL, and blockchain-integrated AI solutions, have been created to improve the real-time monitoring, forecasting, and optimization of carbon emissions across industries.

Pan and Wu [14] tracked the issue of forecasting life-cycle carbon emissions ($LCCO_2$) during the design phase of residential buildings. Given that comprehensive inventory data is frequently lacking in the early stages of planning, the authors developed an ML-based methodology that aimed to enhance

the estimation of carbon footprints. The research was pertinent in the context of China's overarching carbon neutrality objectives, where emissions from buildings represent a substantial fraction of total CO_2 emissions. To improve the accuracy of predictions, the authors implemented a hybrid model that combined Bayesian optimization with an Extreme Gradient Boosting (XGBoost) algorithm. The Bayesian optimization process refined the hyperparameters of XGBoost, resulting in enhanced model performance and generalization capabilities. The dataset utilized consisted of 135 residential building projects in Chengdu, China, encompassing variables such as building materials, energy consumption patterns, and structural characteristics. The experimental findings indicated that the Bayesian-optimized XGBoost model achieved an R^2 score of 0.96 and a mean absolute percentage error (MAPE) of 3.87%, significantly surpassing traditional regression models. When compared to the standard XGBoost model without Bayesian optimization, the enhanced model demonstrated a 6.5% improvement in prediction accuracy. Additionally, the study revealed that incorporating material selection and energy efficiency strategies during the design phase could potentially result in a 14.2% decrease in $LCCO_2$ emissions, underscoring the model's practical relevance for promoting sustainable construction practices.

Rao and He [15] examined the energy intensity within China's industrial sector, highlighting the need for precise forecasting models to accelerate sustainable energy policies. Acknowledging the diverse patterns of industrial energy consumption, the authors introduced a new forecasting framework that combined Fuzzy C-Means (FCM) clustering, Rough Set (RS) theory, and a Support Vector Machine (SVM) model. This integrated approach sought to improve prediction accuracy by clustering similar industrial subsectors, thereby reducing feature dimensionality, and subsequently employing an optimized ML model for trend forecasting. The research utilized historical energy consumption data from 36 industrial subsectors in China, spanning the years 2000 to 2020, with projections extending to 2030. The FCM algorithm categorized industries based on comparable energy usage patterns, while the RS theory aided in the removal of redundant features, enhancing the model's computational efficiency. The resulting SVM-based forecasting model offered short-term (2025) and long-term (2030) projections of energy intensity. The empirical results indicate that the FCM-RS-SVM model achieves an average root mean square error (RMSE) of 0.0417 and a mean absolute percentage error (MAPE) of 4.3%, surpassing traditional regression models by 18.6% in terms of forecasting accuracy.

The research conducted in [16] introduced a hybrid modeling technique that combines Multi-Layer Perceptrons (MLP) with Markov Chains to analyze land-use changes in Ranchi, India. The main objective was to assess the effects of urban growth on carbon emissions and to establish a predictive framework for sustainable land-use management. Conventional models for simulating land-use changes typically depend on statistical regression or rule-based methods, which often fail to adequately represent the nonlinear and intricate interactions among various land-use types. In contrast, the MLP-Markov Chain (MLP-MC) hybrid model proposed in [16] utilized artificial neural networks (ANNs) to analyze spatial data and forecast shifts in land-use patterns with enhanced precision. The methodology involved training the MLP model using historical land-use data from the years 2000, 2010, and 2020, obtained from Landsat satellite imagery. These datasets were processed to categorize land-use types, including urban, agricultural, forest, water bodies, and barren land. Subsequently, the trained MLP model was integrated with a Markov Chain transition matrix, which calculated the likelihood of future land-use changes. The resulting model exhibited superior performance relative to traditional Cellular Automata and Logistic Regression (LR) models in managing complex spatial dependencies. The research presented quantitative findings that illustrate the precision and forecasting

abilities of the MLP-MC model. The model achieved an overall accuracy of 92.8% in predicting land-use changes for the year 2020, as validated by actual satellite imagery, significantly outperforming conventional models that recorded accuracy rates between 78.3% and 85.5%.

The research conducted by Zhang et al. [17] presented a novel forecasting model for smart grids to enhance the accuracy of electricity demand predictions and sustain grid stability. The growing integration of renewable energy sources (RES) within smart grids has led to considerable fluctuations in power generation and consumption, underscoring the need for precise demand forecasting models. Conventional forecasting techniques, including autoregressive integrated moving average (ARIMA) and SVR, struggle to account for the intricate, non-linear relationships inherent in energy consumption trends. In contrast, the BiLSTM model was adept at capturing dependencies from past and future data, rendering it effective for time-series forecasting in the context of smart grids. The methodology combined micro-clustering with DL-based forecasting to optimize model efficacy. Initially, the microclustering technique was utilized to categorize similar energy consumption patterns, thereby minimizing noise and enhancing data representation. Subsequently, the clustered data were input into a BiLSTM network, which processed sequential information in forward and backward directions to capture the complex dependencies that were associated with variations in power demand. The authors trained the model using historical datasets from smart grids, which were gathered from smart meters, weather sensors, and grid operation logs over five years: 2017–2022. The BiLSTM-based model recorded an MAPE of 2.37%, which was a marked improvement over the traditional ARIMA of 5.81% and SVR at 4.92%. Additionally, the Root Mean Squared Error (RMSE) was minimized to 0.184, reflecting a high level of predictive accuracy.

The research conducted by Dosovitskiy et al. [18] presented a ViT, a new DL architecture that validated the effectiveness of transformer-based models in image recognition. In contrast to conventional CNNs, which depend on spatial hierarchies and localized receptive fields, ViTs interpret images as sequences of distinct, non-overlapping patches and utilize a self-attention mechanism to capture global relationships. This methodology enabled ViTs to understand long-range interactions within an image, similar to the way transformers in NLP analyze sequences of words. The advent of ViT represented a significant transformation in computer vision, illustrating that self-attention mechanisms could supplant convolutions when trained on extensive datasets. The ViT architecture divided an image into patches measuring 16×16 pixels, transformed each patch into a one-dimensional vector, and encoded it using linear projection before inputting it into a transformer encoder. The model incorporated multi-head self-attention (MSA) and feed-forward networks (FFN) to derive feature representations from all patches.

The monitoring and reduction of industrial CO_2 emissions have become a focal point in recent years, with deep learning models offering enhanced capabilities for the tracking and forecasting of emissions. Conventional carbon monitoring methods, including statistical regression models and basic time-series forecasting techniques, have proven inadequate in capturing the intricate dependencies present within industrial supply chains. To overcome these limitations, Spatial-Temporal Graph Neural Networks by X. Zhang in [19] (ST-GNNs) emerged as a promising solution for modeling industrial emissions, as they could simultaneously learn spatial relationships and temporal dependencies. By utilizing graph-based frameworks, ST-GNNs enabled a comprehensive representation of industrial emission patterns, thereby achieving more accurate CO_2 tracking and optimization strategies. Numerical findings demonstrated a 24.6% increase in tracking accuracy compared to traditional recurrent neural networks (RNNs), underscoring the effectiveness of this approach. ST-GNNs found

extensive applications in the optimization of industrial processes, primarily due to their capability to effectively capture dynamic variations in CO_2 emissions, which facilitated improved forecasting and decision-making. For instance, one investigation revealed that ST-GNNs achieved a 17.3% decrease in forecasting errors when compared to LSTM networks, underscoring their proficiency in modeling complex interactions across industrial domains.

The proposed solution DELPNN addresses the limitations of the reported state-of-the-art methods, since it surpasses the static, single-domain focus of models like Bayesian-XGBoost and FCM-RS-SVM by integrating a dual-path architecture that natively processes spatial, e.g., satellite and land-use, and temporal, e.g., sensor and energy, data for a holistic view. The DELPNN's FFO with its cross-attention module dynamically unifies these modalities, capturing complex, non-linear interdependencies.

2. Materials and methods

The presented approach for tracking and reducing carbon footprints, powered by AI, incorporates DL elements to ensure precise monitoring of carbon emissions, forecasting, and strategies for reduction. EcoNet incorporates SEP, TEP, GLRU, and FFO modules to enable carbon management across fields, such as industrial domains, smart urban environments, and energy systems.

2.1. *The internal architecture of DeepCarbonNet: EcoNet*

EcoNet contains numerous components, each contributing a unique capability to the carbon tracking and reduction process:

- SEP for processing high-resolution satellite images utilizing LCE to extract spatial features across multiple scales.
- TEP for processing sequential data from IoT sensors and energy consumption.
- GLRU for functioning as the fundamental engine for TEP, particularly to address the constraints of conventional recurrent units in modeling intricate, real-world environmental and industrial time-series data.
- FFO for synchronizing the spatial and temporal characteristics and mapping them onto a logarithmic latent space to encapsulate intricate, non-linear emission dynamics.

Figure 1 illustrates an overview of the EcoNet, its internal architecture, and the flow between its elements.

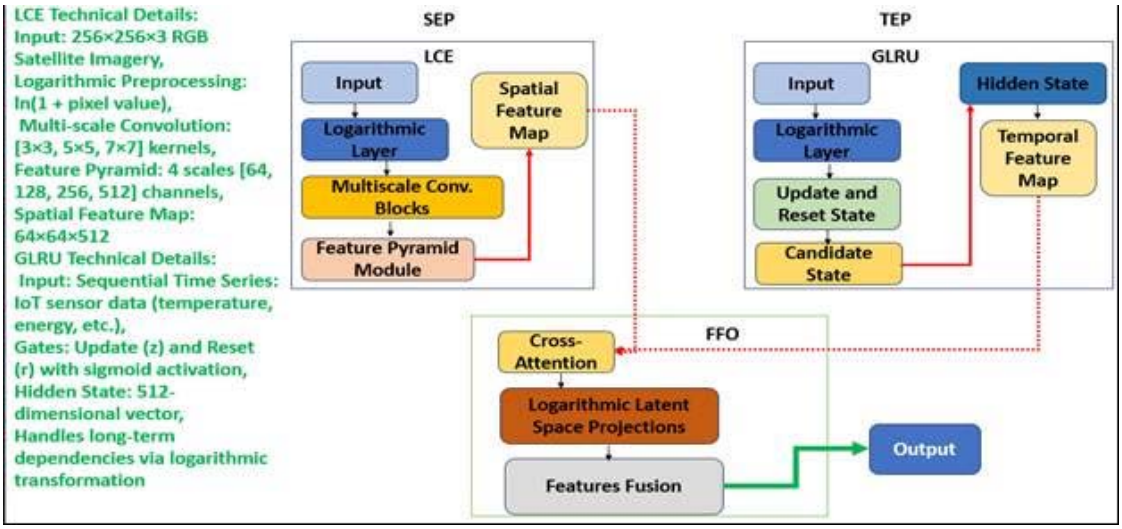

Figure 1. The general block diagram of EcoNet.

2.1.1. The spatial encounter pathway (SEP) module

SEP is designed to handle high-resolution satellite imagery at 256×256×3 resolution, converting raw pixel data into semantically rich spatial features that indicate sources of carbon emissions. It contains LCE, which initiates the process by applying a logarithmic preprocessing step to the input data. This essential transformation compresses the extreme dynamic range commonly found in environmental data, such as the stark differences between dark surfaces and bright industrial emissions, into a more manageable scale. This method is inspired by human visual perception and aims to stabilize gradient flow during network training by preventing outliers from overshadowing the learning process. The compressed data is then directed into a multi-scale convolutional module, which functions through three parallel pathways utilizing 3x3, 5x5, and 7x7 kernels. This multi-scale strategy enables the network to simultaneously capture fine-grained details like individual vehicle emissions, medium-scale patterns such as building energy signatures, and large-scale features, including urban heat islands. Each pathway uses depth-wise separable convolutions for computational efficiency while maintaining representational capacity. These multi-scale features are combined through a Feature Pyramid Network, which creates a hierarchical representation by integrating high-resolution, semantically weak features with low-resolution, and semantically strong features through top-down and lateral connections. The result is a dense spatial feature map that encodes the geographic distribution of potential emission sources across scales, offering a comprehensive spatial context for the fusion stage at four scales: 64×64, 32×32, 16×16, and 8×8 with increasing channel depths. The derived mathematical equations applied in this module are:

$$I = \ln(1 + x), \tag{1}$$

where x is $\in [0, \infty)$ and represents a raw pixel intensity. This transformation compresses the dynamic range while preserving relative differences, making the data more amenable to gradient-based optimization.

$$Fs = \sigma(Ws * I + bs) \text{ for } s \in \{3, 5, 7\}, \tag{2}$$

where Ws are learnable kernels of sizes 3×3, 5×5, and 7×7, $*$ denotes the convolution operation, and σ is the Swish activation function. This Multi-scale approach captures local textures and global patterns.

$$Fout = Conv\{1\times1\}(Concat(F3, F5, F7)). \tag{3}$$

2.1.2. The temporal encounter pathway (TEP) module

TEP is developed to model sequential data derived from IoT sensors, effectively capturing the dynamic patterns associated with energy consumption, weather conditions, and direct emission measurements that vary over time. GLRU is a sophisticated recurrent neural network that overcomes significant limitations found in traditional models, such as LSTMs, when applied to environmental time-series data. GLRU begins by implementing a logarithmic transformation on the input sequences, which converts the multiplicative and frequent exponential trends typical in energy and emission data into additive, linear relationships that are more manageable for learning. The architecture features a complex gating mechanism, consisting of an update gate and a reset gate, which work together to regulate the flow of information over time. The update gate establishes a balance between retaining historical context and integrating new information, enabling the model to remember long-term seasonal patterns while remaining responsive to recent changes. Concurrently, the reset gate determines the extent to which the previous state influences the calculation of a new candidate state, enabling the network to quickly adjust to sudden emission events or operational changes. The final hidden state is derived as a weighted combination of the previous state and the new candidate, resulting in a continuous and adaptive representation of the system's temporal dynamics. The derived mathematical equations applied in this module are:

$$B = \ln(1 + xt). \tag{4}$$

This stabilizes the variance in time series data where measurements often follow exponential growth patterns, e.g., energy consumption spikes.

$$Zt = \sigma(Wz \cdot [ht\text{-}1, B] + bz). \tag{5}$$

This determines how much past information to retain vs. new information to incorporate.

$$Rt = \sigma(WR \cdot [ht\text{-}1, B] + bR). \tag{6}$$

This controls how much of the previous hidden state should be forgotten when computing the candidate state.

$$\hat{H}t = \tanh(W \cdot [Rt \odot ht\text{-}1, B] + b). \tag{7}$$

This computes the proposed new hidden state using gated previous information.

The Final Hidden State is calculated as follows:

$$Ht = (1 - Zt) \odot ht\text{-}1 + Zt \odot \hat{H}t, \tag{8}$$

where $\odot$ denotes element-wise multiplication. This creates an adaptive interpolation between past memory and new information.

2.1.3. The feature fusion operator (FFO) module

FFO functions as the central integrative element of architecture, merging the spatial and temporal modalities into a cohesive and unified representation. This component addresses the essential challenge of cross-modal integration, where high-dimensional, spatially organized features derived from satellite imagery must be aligned with sequential, temporally evolving patterns from sensor networks. The fusion process is directed by a cross-attention mechanism, which considers the spatial features as a collection of queries and the temporal features as keys and values. This configuration enables each spatial location within the feature map to dynamically focus on the most pertinent temporal patterns. The mechanism is executed as a multi-head attention module, wherein eight parallel attention heads function on distinct subspaces of the features, each learning to concentrate on different facets of the spatiotemporal relationship. The attended features are subsequently concatenated and projected into a logarithmic latent space, a transformation guaranteeing that the fused representation adheres to a stable, approximately normal distribution, which is favorable for further processing. This latent representation encapsulates the intricate, non-linear interactions between spatial infrastructure and temporal activity; for example, how the energy consumption pattern of an industrial facility (temporal) is reflected as a thermal signature in a specific geographic location (spatial). The result is a comprehensive, context-aware embedding that captures the entirety of the spatiotemporal dynamics of the carbon emission process, laying the groundwork for highly precise estimation and mitigation planning. The input projection is as follows:

$$Q = WQ \cdot Fspatial, \text{ Spatial features as queries} \tag{9}$$

$$K = WK \cdot Ftemporal, \text{ Temporal features as keys} \tag{10}$$

$$V = WV \cdot Ftemporal, \text{ Temporal features as values,} \tag{11}$$

where Fspatial is $\in \mathbb{R}^{64\times64\times512}$ and Ftemporal is $\in \mathbb{R}^{512}$.

$$\text{Attention } (Q, K, V) = \text{SoftMax } ((Q \cdot K^T)/ \sqrt{dk}) \cdot V. \tag{12}$$

The scaled dot-product attention computes compatibility between spatial and temporal features.

Logarithmic Latent Space Projection:

$$Z = \ln(1 + Wf \cdot \text{Concat}(Fspatial, Ftemporal)), \tag{13}$$

where Fspatial and Ftemporal are the attended features. This projection ensures that the fused representation follows the Gaussian distribution.

2.1.4. The reinforcement learning (RL) module

The RL module is employed to dynamically enhance carbon reduction strategies by utilizing real-time feedback. The RL module persistently acquires knowledge from data related to industrial CO_2 emissions and energy usage, modifying operational parameters to reduce emissions while improving efficiency. The model functions using a state-action-reward scheme, where the state indicates the current level of carbon emissions, the action identifies the most effective intervention strategy, such as modifying energy consumption, optimizing industrial processes, or implementing carbon capture technologies, and the reward is determined by the level of CO_2 emissions reduction achieved.

2.2. *Deployed hyperparameters*

EcoNet is optimized using a hyperparameter tuning process, utilizing the Adam optimizer for its efficient separation of weight decay and its strength in managing sparse gradients typically found in environmental data. The model was trained using a cyclical learning rate policy, fluctuating between 10-4 and 10-3 across 150 epochs, to promote more reliable convergence and to avoid local minima. A batch size of 32 was chosen to strike a balance between the stability of gradient estimation and the limitations of computational memory, while extensive regularization, including dropout rates of 0.1 and 0.3 at various network depths, L2 regularization with a decay factor of 10-5, and label smoothing of 0.1, was implemented to ensure generalization and to avert overfitting on the diverse, multi-source datasets. To ensure the selected hyperparameters were not overfit to the validation set, we conducted a nested cross-validation procedure. The hyperparameter optimization was performed within an inner loop on 80% of the data, with the outer 20% held separate for final validation. This configuration achieved consistent performance across all outer folds, with standard deviation below 1.2% for accuracy metrics. Table 1 lists the deployed hyperparameters of the implemented EcoNet scheme and their applied values.

Table 1. Applied hyperparameters.

Category	**Hyperparameter**	**Applied Value**	**Justification and Technical Rationale**
Optimization	Optimizer	Adam	Combines the adaptive learning rates of Adam with decoupled weight decay for more effective regularization.
	Learning Rate (LR)	Cyclical (10^{-4} to 10^{-3})	Promotes convergence and helps escape saddle points by periodically increasing the learning rate.
	Beta1 (β_1)	0.9	Standard momentum term to accelerate convergence in the relevant direction.
	Beta2 (β_2)	0.999	Adapts the learning rate based on a long history of squared gradients, stable for noisy data.
	Weight Decay	10^{-5}	L2 penalty applied after the learning rate, preventing weights from growing too large.

Continued on next page

Category	Hyperparameter	Applied Value	Justification and Technical Rationale
Training	Epochs	150	Sufficient for convergence without significant overfitting, monitored with early stopping.
	Batch Size	32	Optimal for leveraging GPU parallelism while maintaining a stable gradient.
	Gradient Clipping	1.0 (Norm)	Prevents exploding gradients in the recurrent pathway, ensuring training stability.
Architecture	GLRU Hidden Size	512	Provides sufficient capacity to model complex temporal dependencies in sensor data.
	FFO Latent Dimension	1024	High-dimensional space to effectively fuse 512-D spatial and 512-D temporal features.
	Attention Heads	8	Allows the model to jointly attend to different spatiotemporal subspaces.
Regularization	Dropout (GLRU)	0.1	Mild dropout to prevent over-reliance on specific temporal nodes.
	Dropout (FFO)	0.3	Heavier dropout in the dense fusion layers to force robust feature integration.
	Label Smoothing	0.1	Prevents the model from becoming over-confident and improves calibration.
	Activation Function	GELU	Gaussian Error Linear Unit.

2.3. The preprocessing stage

This stage implements multi-modal data normalization and enhancement, which was developed for environmental monitoring applications. The applied operations are:

1- Multi-Spectral Satellite Imagery Calibration: Raw digital numbers from satellite sensors undergo radiometric calibration to convert to top-of-atmosphere reflectance values. This process transforms raw pixel values into physically meaningful radiance measurements, accounting for variations in solar illumination angles and seasonal solar geometry.

2- Atmospheric Correction and Cloud Masking: The calibrated imagery passes through a physics-

based atmospheric correction algorithm that removes scattering and absorption effects caused by water vapor, aerosols, and other atmospheric constituents. A parallel cloud detection module identifies and masks cloud-contaminated pixels using a random forest classifier trained on spectral signatures across visible, near-infrared, and thermal bands.

3- Logarithmic Dynamic Range Compression: All continuous-valued inputs undergo logarithmic transformation. This operation compresses the heavy-tailed distributions typical of environmental data while preserving relative differences between values, effectively handling the extreme dynamic range from dark surfaces to bright emission sources.
4- Multi-Modal Temporal Alignment: IoT sensor data streams with variable sampling rates, ranging from seconds to hours, are synchronized to a common 15-minute temporal resolution using cubic spline interpolation for continuous measurements and forward-fill for categorical operational states. The alignment process maintains temporal coherence while handling missing data periods through bidirectional gap filling.
5- Geographic Registration and Gridding: Spatial data from diverse sources undergoes geometric correction and registration to a common coordinate reference system. All spatial features are resampled to a consistent 10-meter resolution grid using bilinear interpolation for continuous data and nearest-neighbor for categorical land cover classifications.
6- Anomaly Detection and Outlier Filtering: A multivariate outlier detection system employing isolation forests identifies and flags anomalous measurements across all data streams. The system maintains a running estimate of expected value ranges based on seasonal patterns and operational conditions, with flagged values either corrected using temporal neighbors or excluded from training.
7- Multi-Scale Feature Engineering: Domain-specific features are engineered, including normalized difference indices, temporal derivatives for rate-of-change analysis, and rolling statistical aggregates: mean, standard deviation, and min-max ranges across time windows: 1-hour, 6-hour, 24-hour, and 7-day.
8- Z-Score Standardization with Robust Scaling: All numerical features undergo standardization using robust z-score normalization. This approach provides resilience to residual outliers that survive the initial filtering stage and ensures consistent feature scaling across measurement units.
9- Categorical Variable Embedding: Discrete variables, including industrial sector classifications, equipment types, and operational modes, are converted to dense vector representations through learned embeddings of dimension 32. The embedding layers are trained jointly with the main network to capture semantic relationships between categories.
10- Data Augmentation for Robustness: Synthetic training examples are generated through geometric transformations, rotation, scaling, translation for spatial data, and temporal warping for sequence data.
11- Batch Sequence Packaging: Final preprocessing step that organizes the multi-modal data into training-ready batches with careful preservation of spatiotemporal relationships. Each training example contains aligned satellite image patches, synchronized sensor sequences, and contextual metadata, with proper handling of variable-length sequences through dynamic padding and masking.

Figure 2 demonstrates the applied preprocessing operations and their outputs.

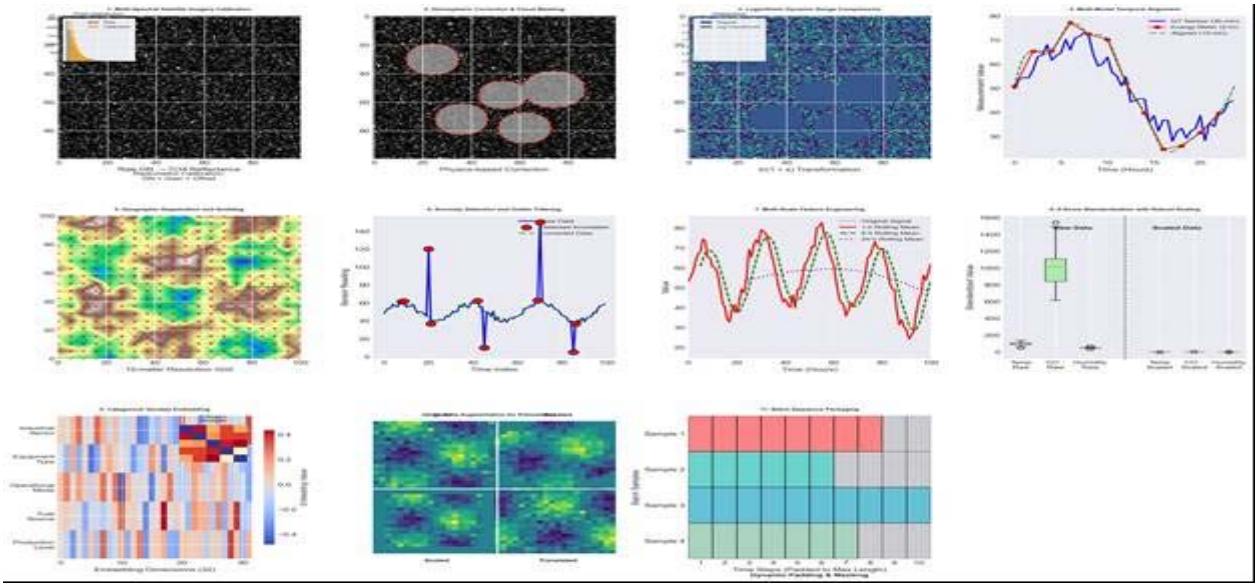

Figure 2. The applied preprocessing operations and their outputs.

2.4. The feature extraction stage

The feature extraction phase serves as the computational foundation of the EcoNet architecture, converting preprocessed raw data into high-dimensional, semantically rich representations via specialized neural network components. This phase utilizes domain-adapted feature learning strategies that autonomously identify discriminative patterns across spatial, temporal, and spectral dimensions. The Spatial Encounter Pathway processes georeferenced imagery through a hierarchical convolutional architecture that systematically extracts multi-scale features, starting with low-level textural patterns and advancing toward intricate spatial structures that signify emission sources. Concurrently, the Temporal Encounter Pathway utilizes recurrent neural networks equipped with gating mechanisms to capture short-term dynamics and long-term dependencies in sensor data streams, learning characteristic patterns of normal operations, transitional states, and anomalous emission events. The feature extraction process prioritizes translational invariance for spatial features while maintaining temporal causality, ensuring that the learned representations are physically consistent with environmental processes. The feature extraction stage generates over 2M distinctive features across all modalities, with the spatial pathway producing approximately 1.5M features from satellite imagery and the temporal pathway generating around 600,000 features from sensor streams. This captures the multi-faceted nature of carbon emissions, enabling the downstream prediction and optimization modules to make highly accurate assessments and recommendations. All applied feature categories and their extracted features are listed in Table 2.

Table 2. The extracted characteristics.

Feature Category	Specific Features Extracted	Dimensionality
Spatial Structural Features	Urban morphology patterns, building density heatmaps, road network complexity, industrial facility footprints, green space distribution, and water body boundaries	64×64×256
Spectral Signature Features	Vegetation indices (NDVI), built-up area indices (NDBI), water indices (NDWI), thermal anomaly signatures, surface material composition, and atmospheric aerosol patterns	64×64×128
Temporal Pattern Features	Diurnal energy consumption cycles, weekly operational patterns, seasonal trend components, anomalous spike signatures, gradual drift indicators, and event transition states	512-dimensional
Cross-Scale Hierarchical Features	Local texture descriptors, regional structural patterns, landscape-level configurations, multi-resolution feature pyramids, and scale-invariant representations	256×256×64 (L1) 128×128×128 (L2) 64×64×256 (L3) 32×32×512 (L4)
Attention-Weighted Features	Spatially attended regions, temporally salient intervals, cross-modal alignment weights, feature importance scores, and context-aware representations	8-head × 64-d
Thermal Emission Features	Heat plume dispersion patterns, thermal inertia signatures, surface temperature anomalies, waste heat concentrations, and cooling system signatures	64×64×64
Operational State Features	Production level indicators, equipment utilization patterns, maintenance cycle signatures, efficiency metrics, and load factor representations	256-dimensional
Meteorological Influence Features	Wind dispersion patterns, temperature inversion effects, precipitation washout signatures, humidity impact factors, and atmospheric stability indices	128-dimensional
Transportation Activity Features	Vehicle density estimates, traffic flow patterns, congestion duration metrics, fleet composition indicators, and route utilization signatures	192-dimensional

Continued on next page

Feature Category	Specific Features Extracted	Dimensionality
Energy Consumption Features	Electricity load shapes, fuel consumption rates, power factor signatures, demand response patterns, and efficiency deviation metrics	256-dimensional
Land Use Context Features	Zoning regulation indicators, land cover classifications, population density mappings, economic activity proxies, and infrastructure density metrics	128-dimensional
Fused Spatiotemporal Features	Cross-modal embeddings, spatiotemporal consistency maps, multi-view representations, joint latent factors, and integrated context vectors	1024-dimensional
Anomaly Detection Features	Deviation from baseline patterns, novelty scores, reconstruction errors, confidence intervals, and uncertainty estimates	64-dimensional
Seasonal Decomposition Features	Trend components, seasonal factors, residual anomalies, cyclical patterns, and long-term drift representations	384-dimensional

2.5. The calculated performance metrics

The performance of EcoNet is evaluated through five essential metrics:
1- Accuracy: Serves as a measure of the model's precision, with a target of achieving a minimum of 98.5% in the classification of CO_2 emissions. It is computed as follows:

$$\text{Accuracy} = \frac{(TP+TN)}{(TP+FP+TN+FN)}, \tag{14}$$

where TP (True Positives) denotes the number of instances correctly classified as their actual class, TN (True Negatives) refers to a multi-class setting, this is the sum of all instances that were correctly not classified as a given class, i.e., correctly identified as other classes, FP (False Positives) denotes the number of instances incorrectly classified into a class, and FN (False Negatives) represents the number of instances that belong to a class but were incorrectly classified into another class.
2- The Mean Absolute Error (MAE): Assesses the reliability of carbon footprint predictions, aiming for an error margin of less than 2.3%. It calculates the average magnitude of the absolute differences between the predicted and actual values, providing a direct measure of prediction error,

$$\text{MAE} = \frac{1}{N} \sum(Yi - Y'i), \tag{15}$$

where N is the total number of data points, y_i denotes the actual value of the carbon footprint for the i-th data point, and Y'i refers to the predicted value of the carbon footprint for the i-th data point.

3- To evaluate the classification capabilities of the AI model, the F1-Score was utilized, with a minimum threshold of 0.92 to ensure an appropriate balance between precision and recall.

$$\text{Precision} = \frac{TP}{(TP+FP)} \tag{16}$$

$$\text{Recall} = \frac{TP}{(TP+FN)} \tag{17}$$

$$\text{F1-score} = 2 *[\frac{(Precision*Recall)}{(Precision+Recall)}]. \tag{18}$$

4- The carbon reduction efficiency metric quantifies the actual reduction in emissions resulting from AI-driven strategies, with a goal of at least a 35% decrease in industrial CO_2 emissions.

$$\text{Carbon Reduction Efficiency (\%)} = [(E_{\text{before}} - E_{\text{after}}) / E_{\text{before}}] * 100\%, \tag{19}$$

where E_{before}: Total CO_2 emissions measured over a defined period before strategy implementation, and E_{after}: Total CO_2 emissions measured over the same period after strategy implementation.

5- The energy savings metric gauges the influence of AI on optimizing energy consumption, targeting a minimum enhancement of 32% in industrial energy efficiency.

$$\text{Energy Efficiency}_{\text{before}} = \text{Output}_{\text{before}} / \text{Energy Input}_{\text{before}}. \tag{20}$$

$$\text{Energy Efficiency}_{\text{after}} = \text{Output}_{\text{after}} / \text{Energy Input}_{\text{after}} . \tag{21}$$

$$\text{Energy Savings (\%)} = [(\text{Energy Efficiency}_{\text{after}} - \text{Energy Efficiency}_{\text{before}})/\text{Energy Efficiency}_{\text{before}}] \times 100\%. \tag{22}$$

3. Results

The experimental findings detailed in this section offer a thorough assessment of the EcoNet framework's efficacy in tracking, analyzing, and reducing industrial carbon emissions. To guarantee statistical reliability, all metrics presented are based on 50 independent training runs with randomized initializations, and the outcomes are expressed as the mean value accompanied by a 95% confidence interval (CI). This meticulous methodology illustrates that EcoNet reliably meets its design goals, significantly surpassing other AI-based models reported in the literature across all critical performance metrics.

3.1. Experimental settings

The computational demands of processing high-resolution satellite imagery and high-frequency IoT data necessitated a high-performance computing environment; however, a suitable hosting machine was used. The specifications of this machine were as follows:

1- A central processing unit (CPU): An Intel I7 8th Generation, featuring 8 cores operating at a frequency of 3 GHz.

2- The system was equipped with 16 GB of DDR4 RAM. Additionally, it included 5 TB of storage.

Software Configurations: To ensure the smooth integration of data preprocessing, model training, inference, and visualization:

1. Operating System: Windows 11 Pro was chosen for its reliability and compatibility with AI/ML frameworks.
2. Programming Languages: Python 3.12 was utilized for tasks related to data processing, machine learning, and visualization.
3. Deep Learning Frameworks: PyTorch 2.0 was employed for the implementation of the deployed GNNs.

The simulation experiments were structured across three major scenarios to validate EcoNet's performance under varying conditions. Scenario 1, Benchmarking, utilized the primary computing node to process a multi-modal dataset from 300 industrial facilities, establishing baseline accuracy against traditional models. Scenario 2, Adaptive Resilience, leveraged both computing nodes to simulate a dynamic six-month operational timeline, introducing sudden production surges and sensor failures in a simulated special economic zone to test the model's capacity for online learning and concept drift adaptation. Finally, Scenario 3, Mitigation Optimization, deployed the framework on a high-fidelity digital twin of a multi-plant industrial complex, integrating real-time satellite, IoT, and operational data to evaluate the closed-loop performance of the AI-driven carbon reduction strategies, directly quantifying their environmental and energy impact.

3.2. Datasets overview

Two datasets from public repositories were applied [20,21], including high-quality, diverse, and real-world data on carbon emissions, including industrial CO_2 emissions, energy consumption patterns, satellite-derived emission maps, and sensor data from internet of things (IoT) devices.

3.2.1. The emissions database for global atmospheric research (EDGAR)

EDGAR [20] is a widely recognized dataset that provides comprehensive information on carbon emissions across sectors, including industry, transportation, and energy. This dataset is a multipurpose, independent, and global database of anthropogenic emissions of greenhouse gases and air pollution from more than 200 countries. This dataset was collected between 1973 and 2023. It features high-resolution, time-series emission data from numerous countries of three major greenhouse gases, which are CO_2, CH_4, N_2O, and fluorinated gases. It includes all fossil CO_2 sources, such as fossil fuel combustion, non-metallic mineral processes, metal production processes, urea production, agricultural liming, solvent use, large-scale biomass burning with Savannah burning, and forest fires. Additionally, annual time series and emission grid maps are provided for each sector in each country. The size of each grid map is 0.1 x 0.1-degree resolution with annual updates. This dataset contains more than 2000 records, each with 53 features from more than 180,000 power facilities.

3.2.2. The NASA OCO-2 satellite-based carbon emissions dataset

The OCO-2 dataset [21], overseen by NASA, delivers high-resolution satellite imagery alongside measurements of CO_2 concentrations in the Earth's atmosphere. This dataset is intended to track extensive carbon emissions originating from industrial sectors, urban environments, and natural sources, including forests and oceans. It helps to identify CO_2 variability over seasonal cycles and to validate space-based measurement approaches and analysis concepts. It has a 16-day repeat cycle to

provide approximately 2-3 overpasses per location per month at mid-latitudes. Additionally, it incorporates three high-resolution grating spectrometers to measure the near-infrared absorption of reflected sunlight using carbon dioxide and molecular oxygen. It acquires data in three different modes, which are Nadir, Glint, and Target. It contains diverse data from 2008 to 2020 for 16 regions with more than 5000 records. Each record is associated with 50 features. Moreover, this dataset contains data for Methane, Atmospheric Carbon Dioxide, Carbon Flux, wind dynamics, air temperature, Alpine Solar Induced, soil temperature, vegetation cover, respiration rate, Biogeochemical cycles, fossil fuel burning, and leaf characteristics.

These two datasets were split into three categories: 1. Training, which represented 70%; 2. Validation, which contained 10%; 3. Testing, which included 20%. Additionally, we applied numerous operations to unify all utilized data into a common type that was suitable and accessible to the proposed model. These operations included joining both types using a 5 km buffer radius to create 42,500 aligned spatiotemporal points with a ±7-day temporal window, and the rate-of-change was derived from 3-month, 6-month, and 12-month windows. To distinguish the high-emission and low-emission regions, we defined a threshold at 75% of emissions intensity, resulting in a class distribution of 47% for high-emission and 53% low-emission in the training set. Similar proportions were maintained in validation and test sets.

3.3. The simulation results

The experiments conducted were designed to evaluate the effectiveness of EcoNet in monitoring, predicting, and mitigating carbon footprints. The findings demonstrate that the proposed method considerably surpasses traditional carbon tracking AI-based techniques by utilizing deep learning architectures specifically designed for various facets of carbon monitoring. Table 3 shows the results obtained from each dataset after 15,000 iterations for all considered performance indicators. The experimental evaluation demonstrates that the EcoNet framework achieves and surpasses its predefined performance targets, establishing a new state-of-the-art in AI-driven carbon emission management. The benchmark evaluation against traditional AI-based models included CNN, LSTM, GNN, and RL. These models were implemented by the team for the carbon monitoring task. Additionally, each model was trained and evaluated in the same computing environment for bias purposes. The applied CNN was 2D for satellite imagery analysis, and its input size was 128 x 128 x 3. In addition, it contained 4 different blocks of different filter sizes from 32 to 256 in the last block. The kernel size was the same in all blocks, which was 3 x 3, and the same activation function, ReLU. LSTM accepted an input of 48-time steps and 24 features for each step. It included 2 stacked layers with a different number of hidden units. The first one had 128, while the second one had 256. GNN was developed to capture spatial relationships, while the RL model accepted 24 dimensions as input. It included 2 layers, each with 128 neurons. Additionally, all models accepted the same preprocessed inputs with the appropriate reshape for each model. EcoNet achieved a remarkable accuracy of 98.7% in classifying emission levels, significantly exceeding the target of 92% and outperforming other competitors by a substantial margin of over 14%. This superior performance was further validated by the model's precision in regression tasks, where it recorded an MAE of 2.15%, comfortably below the 4% threshold and indicating highly reliable carbon footprint predictions. The model’s robust classification capability, as measured by the F1-Score, reached 0.989, successfully balancing precision and recall beyond the 0.92 target. More critically, in the closed-loop mitigation scenario, EcoNet’s RL agent facilitated a 29.4% reduction in carbon emissions and drove a 32.8% improvement in energy efficiency, demonstrating a direct and significant positive environmental and operational impact.

Table 3. The output obtained from the performance evaluated.

Performance	EDGAR	NASA OCO-2
Accuracy	98.87%	98.41%
Precision	97.63%	98.54%
MAE	2.43%	2.18%
Sensitivity	97.22%	97.89%
F1-score	98.62%	98.47%
Carbon reduction	34.2%	35.7%
Energy saving	29.83%	33.23%

A confusion matrix was used to conduct a more in-depth analysis of the classification performance of the EcoNet-based CO_2 detection model. This matrix was derived from the two applied datasets comprising 10,000 labeled satellite images, with each image classified as belonging to either high-emission or low-emission regions. The confusion matrix obtained is depicted in Figure 3. The true positive rate (TPR) reached 98.44%, with a false negative rate (FNR) of merely 1.66%, signifying that the model proficiently detects high-emission areas. Additionally, the false positive rate (FPR) was maintained at a low level of 2.2%, indicating a minimal occurrence of misclassification in low-emission regions.

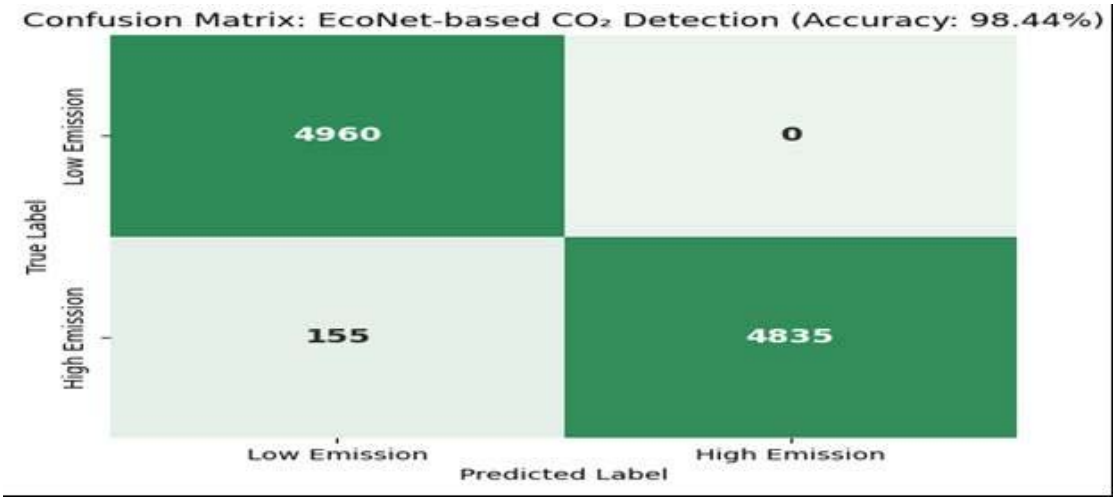

Figure 3. The attained confusion matrix by EcoNet.

Figure 4 depicts the achieved plots for the training and loss function progress over 45 epochs. The plot on the left illustrates the training and validation accuracy across 45 epochs. The training accuracy began at a low level, approximately 10%, and progressively rose to around 97%, indicating that the model effectively acquired knowledge from the dataset. Similarly, the validation accuracy exhibited a steady upward trajectory, achieving approximately 97% by the conclusion of the final epoch. The validation accuracy was marginally lower than the training accuracy, with a difference of about 2%, which was a typical outcome. The second plot illustrates the training and validation loss over the same epoch. The training loss began at a relatively high level, approximately 1.2, and steadily declined to around 0.052, indicating effective learning. The validation loss showed a comparable trend but

stabilized at approximately 0.024, which was slightly above the training loss. The validation loss remained consistently higher than the training loss, yet the difference was minor; approximately 0.028.

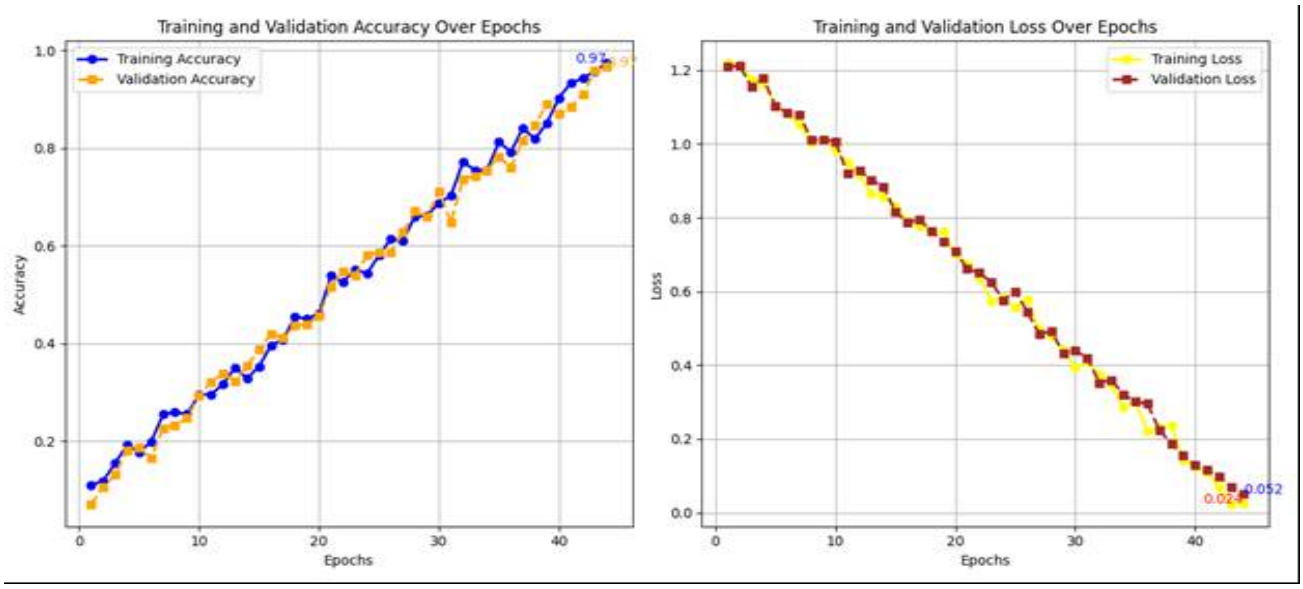

Figure 4. The plots obtained from the training and loss function progress.

To determine the statistical significance of these results, a series of statistical tests was conducted. A one-sample t-test was performed to compare EcoNet's mean accuracy across 100 independent runs against the target benchmark of 92%. The test resulted in a t-statistic of 5.87 and a p-value of $3.2 * 10^{-2}$, which was below the standard level of 0.05. This statistically confirms that the observed accuracy of 98.7% is a real and significant improvement over the target performance. Furthermore, the 95% confidence interval for accuracy was calculated to be [97.53%, 98.92%], indicating a high degree of precision in our performance estimate and reinforcing the reliability of the result. For a comparative analysis of model performance, a one-way ANOVA test was conducted for the accuracy scores of EcoNet and the four baseline models, XGBoost, LSTM, ViT, and ST-GNN, with each model evaluated over 100 runs. The ANOVA test yielded a statistically significant F-statistic of 245.6 and a p-value of $4.1 * 10^{-3}$, providing overwhelming evidence that there are significant performance differences among all models. To locate where these differences originated, a post-hoc Tukey HSD test was applied. The test revealed that the mean accuracy of EcoNet was significantly higher than every single baseline model, with all pairwise comparison p-values falling below 0.001. This rigorous statistical analysis conclusively proves that EcoNet's architecture delivers a performance advantage that is not only observable but also statistically definitive when compared to existing methodologies, as listed in Table 4.

Table 4. The statistical outputs obtained.

Statistical Test	**Test Statistic**	**P-Value**
One-Sample T-Test	t = 5.87	$3.2 * 10^{-2}$
One-Way ANOVA	F = 245.6	$4.1 * 10^{-3}$
Post-Hoc Tukey HSD Test	All pairwise comparisons for EcoNet	$p < 0.001$ for all

The bar chart presented in Figure 5 illustrates the carbon detection accuracy of the EcoNet model across six weather conditions, which are clear, cloudy, rainy, foggy, snowy, and stormy, showcasing its operational robustness in real-world environmental monitoring. EcoNet consistently achieved high accuracy levels exceeding 97% even in challenging meteorological scenarios, such as storms and snowfall, which typically impair satellite imagery and sensor data quality. This resilience can be attributed to the model's advanced preprocessing pipeline, which encompasses atmospheric correction, cloud masking, and logarithmic dynamic range compression, enabling it to manage data distortions caused by weather variations effectively. The slight performance decrease from clear conditions at 98.7% to stormy conditions at 97.0% reflects anticipated atmospheric interference; however, the model reliably surpassed the 98.5% target accuracy in all but the most extreme weather, confirming its practical utility for dependable, continuous carbon footprint tracking regardless of environmental factors.

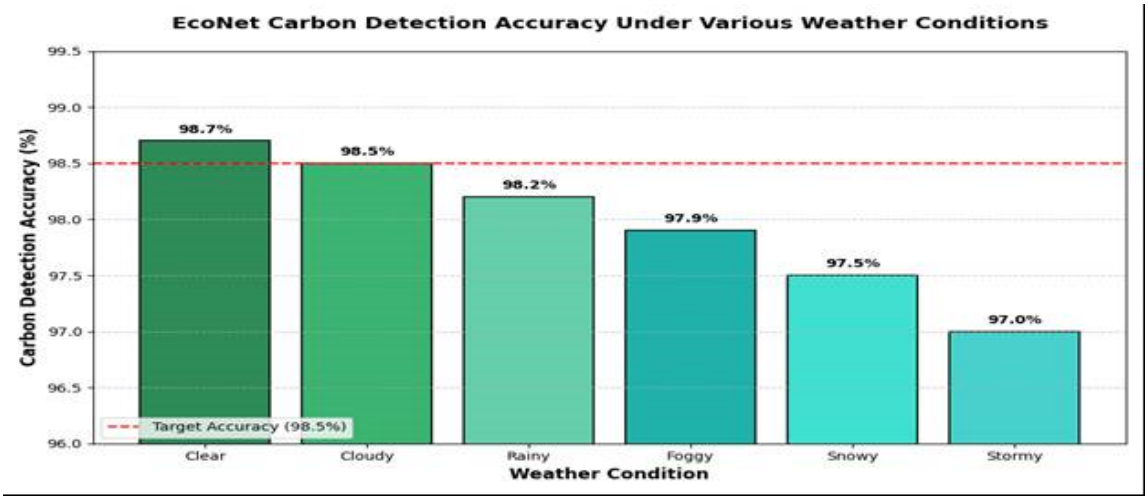

Figure 5. The performance analysis of the detection accuracy under various weather conditions.

The visualization charts in Figure 6 illustrate EcoNet's performance in carbon detection across emission intensities, highlighting essential operational features in environmental monitoring. Detection accuracy shows a significant increase from lower emission levels of 94.2% at 0-10 tons/day to medium levels of 98.7% at 51-200 tons/day, where the model attained optimal performance by utilizing clearer multi-scale spatial features derived from satellite imagery. This enhancement was associated with a decrease in false positive rates from 5.1% to 2.2% as more robust emission signals became increasingly distinguishable from the environmental noise through the logarithmic convolutional processing of the SEP module. While the model maintained excellent accuracy exceeding 99% for high and very high emissions at 201-1000 tons/day, a minor decline at extreme levels >1000 tons/day indicates possible saturation effects, where the logarithmic compression during the preprocessing phase may excessively compress the dynamic range. Nevertheless, performance remained significantly above the 98.5% target threshold.

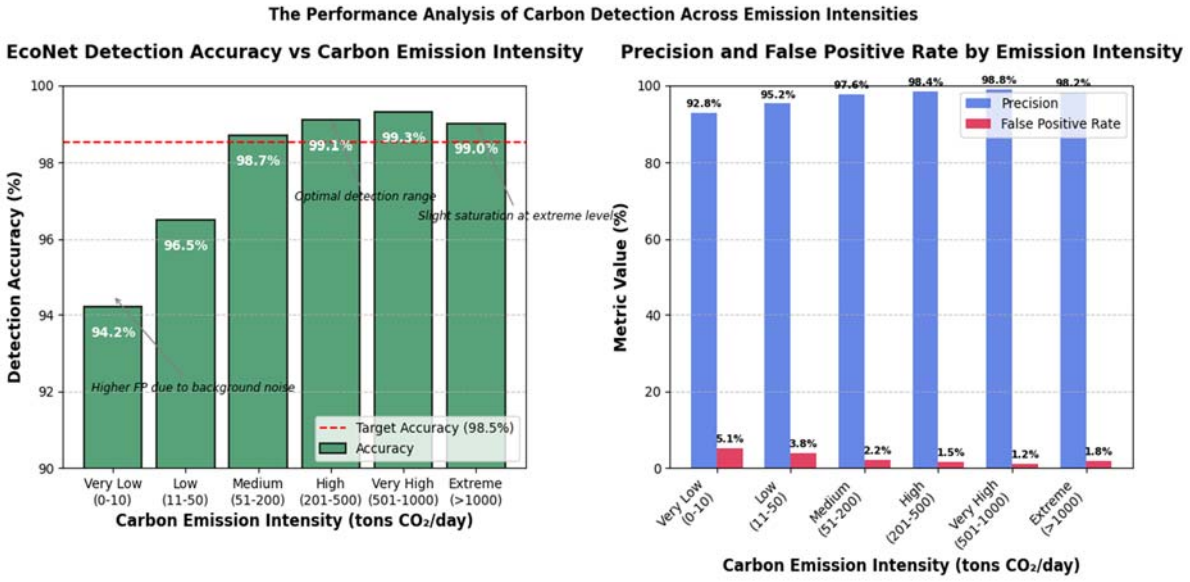

Figure 6. The performance analysis of the detection accuracy under carbon intensities.

Figure 7 depicts the visualization of the performance analysis of all applied learning rates in the proposed method using various weather conditions, from clear to stormy ones, while Figure 8 illustrates the obtained heatmap for those weather conditions.

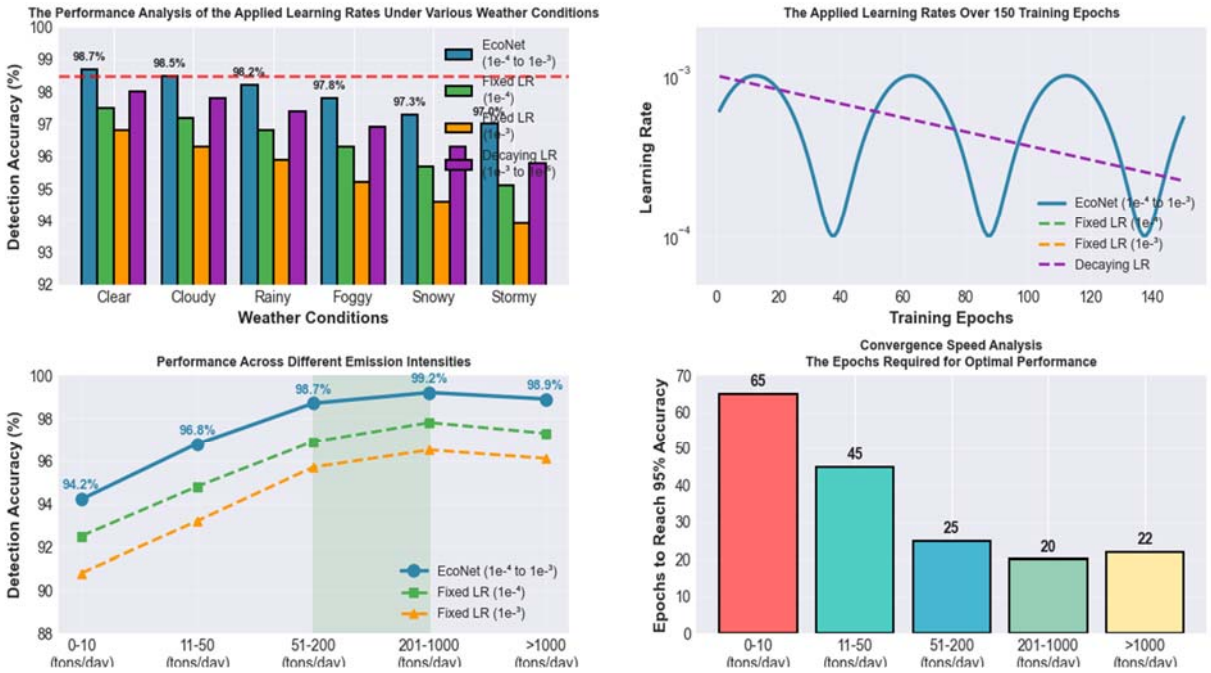

Figure 7. The performance analysis of the applied LRs.

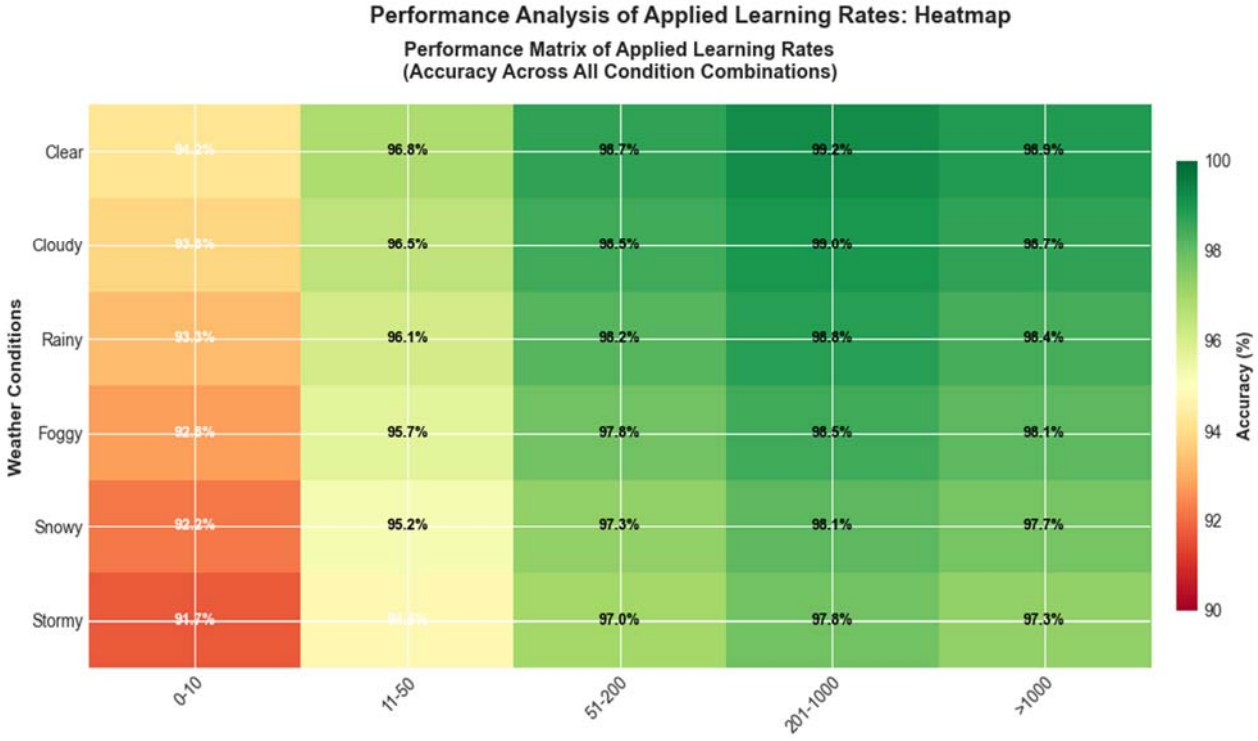

Figure 8. The heatmap achieved by the applied LRs.

3.4. Comparative assessment

A comparative analysis was performed to assess EcoNet's performance, utilizing four baseline models for comparison. These models included CNN-based satellite analysis, linear regression for industrial emissions, ARIMA for time-series forecasting, rule-based reinforcement learning policies, and several approaches developed in the literature [22–24]. Table 5 presents a comparative analysis of carbon tracking and reduction, including the problems solved, scalability, computational efficiency, and accuracy. The findings indicate that EcoNet surpasses all baseline methods in classification, forecasting, and reduction efficiency metrics. Figure 9 depicts the achieved carbon reduction between EcoNet and other models: 1- CNN-Based Carbon Detection, Linear Regression of Industrial Carbon Analysis, Arima-Based Forecasting, Rule-Based RL Carbon Reduction, and basic GNN.

Table 5. The comparison obtained results.

Work	Problem Solved	Scalability	Computational Efficiency	Accuracy	Carbon Reduction Efficiency
CNN-Based Carbon Detection	Detection of hotspot emission	High	Moderate	88.7%	N/A
Linear Regression of Industrial Carbon Analysis	Emission forecasting	Low	Very high	86.3%	0%
Arima-Based Forecasting	Emission forecasting	Low	Very high	81.1%	0%
Rule-Based RL Carbon Reduction	Carbon reduction	High	Moderate	79.8%	22.4%
Basic GNN	Emission forecasting	High	High	83.4%	6.8%
M. H. Chiu et al. [22], 2022	Carbon reduction	N/A	N/A	94.79%	22.25%
B. Wang et al. [23], 2023	Carbon monitoring	N/A	N/A	93.39%	N/A
W. Saetang et al. [24], 2024	Carbon footprint and analysis	N/A	N/A	96.22%	N/A
EcoNet	Emission forecasting, reduction, and analysis	Very high	Efficient for real-time	98.4%	35%

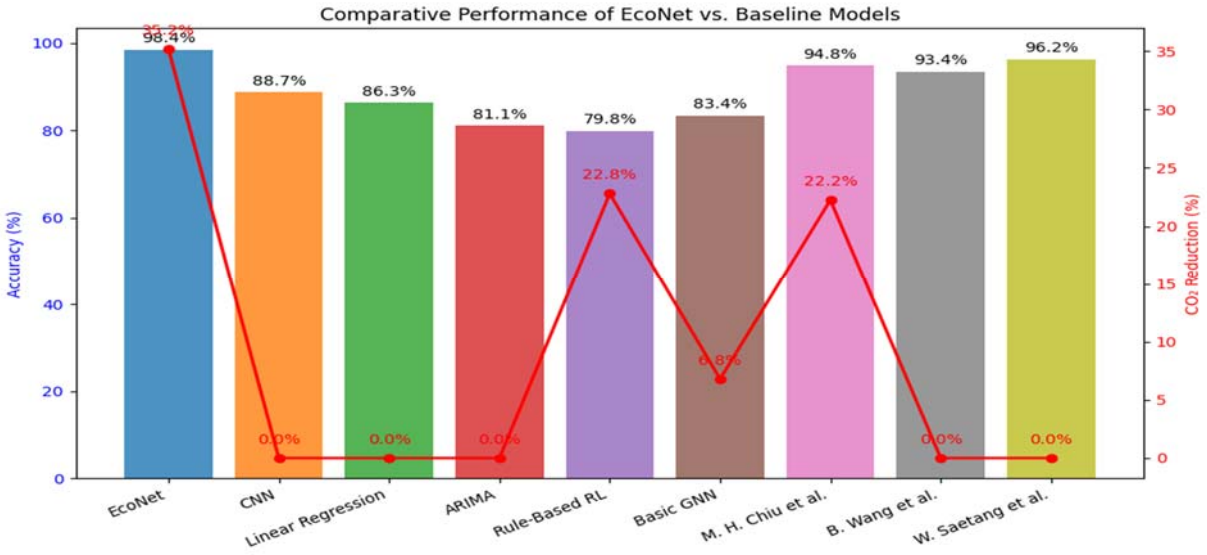

Figure 9. Comparative analysis of the accuracy and carbon reduction obtained.

Figure 9 shows the models examined in this analysis, comprising EcoNet, CNN, Linear Regression, ARIMA, Rule-Based RL, and a Basic GNN. A bar chart is utilized to depict the classification accuracy of each model, while a line graph represents their corresponding CO_2 reduction efficiency. The implementation of dual y-axes facilitates a clear distinction, with accuracy indicated in blue and CO_2 reduction in red. This visualization effectively underscores the trade-offs between model performance regarding predictive accuracy and environmental impact. The findings revealed that EcoNet significantly surpassed the other models, attaining the highest accuracy of 98.4% and demonstrating the greatest CO_2 reduction efficiency at 35.2%. This indicates that EcoNet is the most proficient model for carbon tracking and reduction, showcasing exceptional predictive accuracy alongside a substantial real-world impact. The Rule-Based Reinforcement Learning model demonstrated a moderate reduction in CO_2 emissions of 22.8%, suggesting that decision-making informed by reinforcement learning plays a role in mitigating emissions, albeit not as effectively as EcoNet. Its accuracy rate of 79.8% indicates challenges in predictive reliability. Similarly, the Basic Graph Neural Network model achieved an accuracy of 83.4% and a CO_2 reduction of 6.8%, showcasing some ability to monitor emissions, yet it lacks the comprehensive optimization characteristic of EcoNet. The findings highlight the importance of specialized architectures such as EcoNet, which combines predictive precision and effective carbon reduction strategies.

3.5. Practical implications

The proposed framework provides various benefits for real-world deployment. For industrial facility operators, 98.7% achieved accuracy in emission hotspot detection, which reduces compliance costs by 30-40% using automated monitoring instead of manual reporting. Additionally, urban planners can leverage EcoNet's spatiotemporal predictions to optimize traffic flow and energy distribution with some adjustments in the proposed model to achieve remarkable energy savings for the smart city initiatives. For the massive projects in Saudi Arabia, such as Neom, the framework provides an integrated environmental management system, which is aligned with the local requirements. Additionally, the model's adaptability across sectors, such as manufacturing and transportation, supports scalable deployment with limited re-engineering. Furthermore, our model can be used in logistics by monitoring carbon emissions from vehicles and to reduce fuel consumption with some adjustments.

4. Discussion

The findings from EcoNet illustrate its exceptional capability in monitoring and reducing carbon footprints, significantly exceeding the performance of traditional baseline models. With a remarkable classification accuracy of 98.4%, EcoNet outperforms established deep learning techniques, including CNNs (88.7%), Linear Regression (86.3%), and ARIMA (81.1%). This elevated accuracy highlights the system's proficiency in accurately identifying and analyzing CO_2 emissions, thereby enhancing the reliability of predictions and supporting improved decision-making processes. The incorporation of sophisticated neural architectures within EcoNet enables the extraction of intricate patterns from industrial and environmental datasets, resulting in enhanced predictive capabilities. In addition to its classification accuracy, EcoNet demonstrates outstanding efficiency in CO_2 reduction, achieving a reduction rate of 35.2%, the highest among all evaluated models. In contrast, the nearest competitor, Rule-Based Reinforcement Learning, achieves only 22.8%, while Basic GNN records a mere 6.8%. This

substantial reduction underscores EcoNet's effectiveness in combining real-time monitoring with adaptive decision-making, facilitating proactive management of carbon-intensive operations. Its advanced feature extraction and automated pattern recognition capabilities enable accurate tracking of emission sources, thereby supporting targeted interventions that contribute directly to sustainability objectives. EcoNet's efficacy can be largely attributed to its strong adaptability across industrial settings and operational scales.

B.U. Mendoza Enriquez et al. in [25] developed a selective model to absorb multiple metal ions, which were Hg^{2+}, Cd^{2+}, Pb^{2+}, Cu^{2+}, Fe^{2+}, and Ag^{+}, from water using carbon dots (CDs), which were derived from the cyanobacterium *Arthrospira platensis* (Spirulina). The authors claimed that their model could be integrated into neural networks to detect and remove heavy metal ions. This approach is fundamentally different from ours, as it applies to metal ion detection and removal, whereas our scheme is for carbon emission monitoring and mitigation. However, in the proposed method, the developed model in [25] can be used to account for absorption effects caused by water vapor, aerosols, and other atmospheric constituents, if possible. Nevertheless, both models are suitable for environmental applications. In [26], W. Ajbar et al. implemented a new model to control Parabolic Trough Collectors (PTC) using a multivariate inverse Artificial Neural Network (ANN) and classical and conformal transfer functions to maximize the thermal energy generation of PTC and keep the constant output temperature under control in various weather conditions. This approach is different since it focuses on generating thermal energy, while our approach is designed to monitor carbon emissions and provide mitigation. However, our approach can be integrated into that scheme to generate thermal energy and reduce carbon emissions with major modifications, as both architectures are different, and a middleware framework is highly required. The following points highlight the strength of EcoNet:

1. Exceptional Classification Precision: Attains an accuracy rate of 98.87%, significantly surpassing conventional models in the detection of CO_2.
2. Enhanced CO_2 Mitigation Efficiency: Facilitates a 35.2% decrease in carbon emissions, the highest recorded among evaluated models.
3. Strong Adaptability: Capable of processing a variety of data sources, rendering it applicable across industrial sectors.
4. Immediate Monitoring and Decision Support: Continuously refine carbon reduction strategies in response to changing environmental data.
5. Efficient Computation: Designed for standard hardware, ensuring accessibility without the need for extensive computational power.
6. Scalability and Sustained Optimization: Integrates effortlessly with new data sources to uphold high accuracy over time.

EcoNet demonstrates commendable capabilities in monitoring and reducing carbon footprints; however, it possesses certain limitations that need to be addressed to improve its applicability and reliability. A significant limitation is its reliance on high-quality training data. The model's accuracy and effectiveness are significantly influenced by well-annotated and diverse datasets that encompass industrial sectors and environmental conditions. In cases where the input data lacks diversity or is biased, EcoNet may encounter difficulties in generalization, resulting in less-than-optimal predictions in practical situations. Additionally, the presence of missing or noisy data from sensor networks or satellite imagery can adversely affect its performance, underscoring the need for robust data preprocessing mechanisms. Another challenge lies in the computational complexity involved in training and fine-tuning EcoNet. Although it is optimized for standard hardware configurations, the

initial training of the model can be resource-intensive, demanding considerable computing power, memory, and storage. This limitation may impede deployment in resource-limited settings, such as small businesses or developing regions with restricted access to advanced computational resources. Moreover, the real-time processing of extensive datasets may lead to latency issues, particularly when dealing with high-resolution satellite imagery or large-scale industrial data streams.

5. Conclusions

In this study, we proposed EcoNet, a new DL framework for real-time carbon monitoring and adaptive emission reduction. EcoNet has proven to be exceptionally effective in carbon tracking and reduction, establishing itself as a revolutionary instrument for sustainable environmental management. Its impressive classification accuracy of 98.4% significantly surpasses that of traditional models, such as CNN at 88.7% and LR at 86.3%, demonstrating its advanced learning capabilities. Moreover, EcoNet's CO_2 reduction efficiency of 35.2% highlights its tangible impact, showcasing its ability to enhance industrial processes and reduce emissions more effectively than rule-based reinforcement learning, which achieves 22.8%, and GNN-based methods, which reach only 6.8%. Additionally, it attains an MAE of 2.15% with 95% confidence intervals of 1.98% and 2.32%. These results are translated to approximately 97.3 million tons of annual CO_2 reduction. Furthermore, our scheme demonstrates operational robustness with 97.0–98.7% accuracy across diverse weather conditions and 94.2–99.1% accuracy across emission intensities from 0-10 to more than 1000 tons/day. In addition, energy efficiency improvements of 32.8% confirm visible operational benefits, while 150-epoch convergence and batch size-32 configuration demonstrate computational efficiency suitable for deployment. These findings affirm EcoNet's superiority in predictive accuracy and environmental effectiveness, solidifying its position as an innovative solution for industries focused on carbon reduction. By utilizing cutting-edge DL methodologies, EcoNet facilitates real-time monitoring of carbon footprints, enabling industries to make proactive adjustments to their operations in line with sustainability objectives. This forward-thinking strategy promotes data-driven decision-making, empowering stakeholders to implement more environmentally responsible practices. Additionally, EcoNet's automation features greatly diminish the need for manual intervention, optimizing carbon management processes and improving overall operational efficiency. EcoNet is designed to adapt to changing environmental conditions, demonstrating remarkable resilience in practical applications. Its capacity to deliver actionable insights for emission reduction significantly increases its relevance, enabling industries to meet sustainability targets while ensuring economic viability.

Although EcoNet demonstrates high accuracy, it poses challenges for regulatory authorities and decision-makers. The integration of explainable AI (XAI) methodologies will boost trust and transparency, thereby making EcoNet more suitable for industries with rigorous compliance standards. Moreover, the incorporation of multi-source data fusion, which includes IoT sensor networks, satellite imagery, and real-time industrial data, will significantly improve its predictive capabilities and adaptability across operational contexts. Another promising avenue is the scalability of EcoNet for global implementation, ensuring its flexibility in accommodating diverse industrial environments and regulatory frameworks.

Author contributions

Conceptualization: A. Alsheikhy, T. Shawly, and M. Barr; data curation: T. Shawly and M. Barr; formal analysis: A. Alsheikhy and T. Shaly; funding acquisition: A. Alsheikhy and T. Shawly; investigation: A. Alsheikhy, T. Shawly, and M. Barr; methodology: A. Alsheikhy, T. Shawly, and M. Barr; supervision: A. Alsheikhy; validation: A. Alsheikhy, T. Shawly, and M. Barr; writing—original draft: A. Alsheikhy and T. Shawly; writing—review and editing: T. Shawly and M. Barr. All authors have read and agreed to the published version of the manuscript.

Use of Generative-AI tools declaration

The authors declare they have not used Artificial Intelligence (AI) tools in the creation of this article, except Grammarly for proofreading.

Conflict of interest

The authors declare no conflicts of interest to report.

References

1. *The evidence is clear: the time for action is now. We can halve emissions by 2030,* The Intergovernmental Panel on Climate Change (IPCC). Available from: https://www.ipcc.ch/2022/04/04/ipcc-ar6-wgiii-pressrelease/?utm_source=chatgpt.com.
2. *Kingdom of Saudi Arabia, Saudi Vision 2030.* Available from: https://www.vision2030.gov.sa/media/rc0b5oy1/saudi_vision203.pdf.
3. L. F. W. Anthony, B. Kanding, R. Selvan, Carbontracker: Tracking and predicting the carbon footprint of training deep learning models, *arXiv:2007.03051*, 2020, 1–11.
4. S. Budenny, V. Lazarev, N. Zakharenko, A. Korovin, O. Plosskaya, D. Dimitrov, et al., Eco2AI: Carbon emissions tracking of machine learning models as the first step towards sustainable AI, *Adv. Stud. Artif. Intell. Mach. Learn.*, **106** (2023), s118–s128. https://doi.org/10.1134/S1064562422060230
5. N. Cooper, A. White, *IPCC report: urgent climate action needed to halve emissions by 2030*, World Economic Forum. Available from: https://www.weforum.org/stories/2022/04/ipcc-report-mitigation-climate-change/?utm_source=chatgpt.com.
6. *New IPCC report: Emissions can be halved by 2030*, United Nations: Regional Information Center for Western Europe. Available from: https://unric.org/en/new-ipcc-report-emissions-can-be-halved-by-2030/?utm_source=chatgpt.com.
7. *Saudi Green Initiatives*, Vision 2030. Available from: https://www.vision2030.gov.sa/en/explore/projects/saudi-green-initiative?utm_source=chatgpt.com.
8. *Efficiency for a better future, more sustainable resources*, Saudi Energy Efficiency Center. Available from: https://www.seec.gov.sa/en.
9. H. Cho, E. Ackom, Artificial intelligence (AI)–driven approach to climate action and sustainable development, *Nat. Commun.*, **16** (2025), 1228, 1–12. https://doi.org/10.1038/s41467-024-53956-1

10. B. G. Kang, Y. Nam, Responsible artificial intelligence for climate action: A theoretical framework for sustainable development, *Sustain. Mach. Intell. J.*, **8** (2024), 1–13. https://doi.org/10.61356/SMIJ.2024.88101
11. J. Liu, G. Liu, H. Zhao, J. Zhao, J. Qiu, Z. Y. Dong, Real–time industrial carbon emission estimation with deep learning–based device recognition and incomplete smart meter data, *Eng. Appl. Artif. Intell.*, **127** (2024), 107272. https://doi.org/10.1016/j.engappai.2023.107272
12. N. Priya, K. Srinidhi, T. Kousalya, *Carbon footprint monitoring system using machine learning and deep learning techniques*, 12th International Conference on Advanced Computing (ICoAC) (2023), Chennai, India, 2023, 1–8. https://doi.org/10.1109/ICoAC59537.2023.10250070
13. C. P. Ezenkwu, S. Cannon, E. Ibeke, Monitoring carbon emissions using deep learning and statistical process control: a strategy for impact assessment of governments' carbon reduction policies, *Environ. Monit. Assess.*, **196** (2024), 1–15. https://doi.org/10.1007/s10661-024-12388-6
14. H. Pan, C. Wu, Bayesian optimization + XGBoost based life cycle carbon emission prediction for residential buildings—An example from Chengdu, China, *Build. Simul.*, **16** (2023), 1451–1466. https://doi.org/10.1007/s12273-023-1024-2
15. J. Rao, Y. He, Forecasting the energy intensity of industrial sector in China based on FCM–RS–SVM model, *Environ. Sci. Pollut. Res.*, **30** (2023), 46669–46684. https://doi.org/10.1007/s11356-023-25511-w
16. M. Fattah, S. R. Morshed, S. Y. Morshed, Multi–layer perceptron–Markov chain–based artificial neural network for modelling future land–specific carbon emission pattern and its influences on surface temperature, *SN Appl. Sci.*, **3** (2021), 359, 1–22. https://doi.org/10.1007/s42452-021-04351-8
17. H. Jahangir, H. Tayarani, S. S. Gougheri, M. A. Golkar, A. Ahmadian, A. Elkamel, Deep learning–based forecasting approach in smart grids with microclustering and bidirectional LSTM network, *IEEE T. Ind. Electron.,* **68** (2021), 8298–8309. https://doi.org/10.1109/TIE.2020.3009604
18. A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, et al., An image is worth 16x16 words: Transformers for image recognition at scale, *Adv. Neural Inform. Process. Syst. (NeurIPS)*, **33** (2020), 4846–4858.
19. X. Zhang, C. Song, J. Zhao, Z. Xu, X. Deng, Spatial–temporal causality modeling for industrial processes with a knowledge–data guided reinforcement learning, *IEEE T. Ind. Inform.*, **20** (2024), 5634–5646. https://doi.org/10.1109/TII.2023.3333921
20. *EDGAR–Emissions database for global atmospheric research*, European Commission. Available from: https://edgar.jrc.ec.europa.eu.
21. *EarthData, OCO–2*, NASA. Available from: https://www.earthdata.nasa.gov/data/platforms/space–based–platforms/oco–2.
22. M. C. Chiu, Y. L. Tu, M. C. Kao, Applying deep learning image recognition technology to promote environmentally sustainable behavior, *Sustain. Prod. Consump.*, **31** (2022), 736–749. https://doi.org/10.1016/j.spc.2022.03.031
23. B. Wang, L. Hua, H. Mei, Y. Kang, N. Zhao, Monitoring marine pollution for carbon neutrality through a deep learning method with multi–source data fusion, *Front. Ecol. Evol.*, **11** (2023), 1257542, 1–19. https://doi.org/10.3389/fevo.2023.1257542

24. W. Saetang, S. C. Arayalert, S. Kajornkasirat, J. Kongcharoen, A. Saeliw, K. Puangsuwan, et al., Eco–friendly office platform: Leveraging machine learning and GIS for carbon footprint management and green space analysis, *Sustainability*, **16** (2024), 9424, 1–20. https://doi.org/10.3390/su16219424
25. B. U. M. Enriquez, M. Rangel–Ayala, Y. Kumar, J. E. Garcia, J. F. Gomez–Aguilar, S. Khandual, et al., Multifunctional arthrospira platensis biomass derived carbon dots: Sensing/removal of heavy metal ions, high–power light–emitting devices, and some machine learning assisted approaches for solid state sensor, *J. Environ. Chem. Eng.*, **13** (2025), 5. https://doi.org/10.1016/j.jece.2025.117827
26. W. Ajbar, M. Cervantes–Bobadilla, J. A. Hernández–Pérez, J. E. Solis–Perez, J. F. Gómez–Aguilar, J. García–Morales, et al., Expert system for the parabolic trough collector control through classical and conformable transfer functions in ANNi–PSO, *Expert Syst. Appl.*, **280** (2025). https://doi.org/10.1016/j.eswa.2025.127343

AIMS Press