ISSN 2277-2685
IJESR/April-June. 2026/ Vol-16/Issue-2/1154-1160
Sadiya Fatima *et. al.*, / International Journal of Engineering & Science Research

Full Length Article

Machine Learning Approaches For Real-Time Carbon Emission Prediction And Mitigation

Sadiya Fatima[1], Mohammed Mubashir Ahmed[2], Salwa Baig[3], Mr. Mohammed Mateen Ahmmed[4]
[1,2,3]B.E. Students; Department of IT, Lords Institute of Engineering and Technology, Hyderabad, India.
[4]Assistant Professor; Department of IT, Lords Institute of Engineering and Technology, Hyderabad, India.
Mail Id; mateen.mohd@lords.ac.in[4]

Accepted 30-04-2026

Abstract

The project, "CO2 Emission Rating by Vehicles Using Data Science," is a data-driven initiative aimed at assessing and rating the carbon dioxide (CO2) emissions of new light-duty vehicles available for retail sale in Canada in 2022. Leveraging the power of Python programming and employing sophisticated machine learning models, like the Random Forest Classifier and the Decision Tree Classifier, this project offers a comprehensive analysis of vehicle emissions. The dataset utilized for this project contains crucial information, including fuel consumption ratings, CO2 emissions in grams per kilometer, CO2 ratings on a scale from 1 (worst) to 10 (best), and smog ratings on a scale from 1 (worst) to 10 (best). These data elements provide a holistic perspective on the environmental performance of various vehicle models, allowing consumers and policymakers to make informed choices. We are planning to employ the Random Forest Classifier, a powerful ensemble learning algorithm, and the Decision Tree Classifier to build predictive models to achieve better training and testing accuracy by combining these advanced algorithms and a rich dataset, this project aims to contribute to sustainable transportation solutions and empowers consumers to make environmentally conscious decisions when purchasing vehicles. The CO2 Emission Rating system developed here serves as a valuable tool for evaluating the environmental impact of different vehicle models, helping reduce carbon emissions and mitigate climate change. The developed linear regression model achieved an accuracy of 99.447%, demonstrating its robustness in predicting vehicle CO_2 emissions and supporting reliable real-time mitigation strategies.

Keywords: *Carbon Emission Rating, Machine Learning Models, LSTM, Fuel Consumption, Decision Tree Classifier, Random Forest Classifier, Linear Regression Accuracy, Deep Learning, Ensemble Methods, Data Science For Sustainability.*

Introduction

In the epoch of the Anthropocene, where human activities have become the dominant force shaping the Earth's climate and ecosystems, the surge in carbon dioxide (CO_2) emissions stands as a stark testament to our industrial and technological pursuits. From the relentless combustion of fossil fuels to the rapid expansion of urban infrastructure, the imprint of human progress is now inextricably linked with environmental degradation. The consequences of these emissions are not abstract they are tangible and increasingly severe: rising global temperatures, melting polar ice caps, ocean acidification, and an uptick in the frequency and intensity of extreme weather events such as hurricanes, droughts, and wildfires. These phenomena underscore the urgency with which we must address the climate crisis[1].

As the global community grapples with the multifaceted challenges posed by climate change, a nuanced understanding of CO_2 emissions and their drivers becomes not just important, but imperative. Historically, nations and environmental agencies have relied on traditional statistical methods, manual data collection, and retrospective analyses to estimate and monitor their carbon footprints. While these approaches have laid the groundwork for climate policy and awareness, they often fall short in capturing the dynamic, interconnected, and rapidly evolving nature of modern economies and ecosystems. The complexity of emission sources ranging from industrial processes and transportation networks to agricultural practices and energy consumption demands a more agile, intelligent, and predictive framework[2].

Enter the realm of machine learning (ML) and deep learning (DL) subfields of artificial intelligence that have already transformed industries such as healthcare, finance, marketing, and cybersecurity. Their application in environmental science, particularly in the prediction and mitigation of CO_2 emissions, remains a relatively nascent but profoundly promising frontier. These technologies offer the ability to process vast volumes of heterogeneous data, uncover hidden patterns, and generate predictive insights that are beyond the reach of conventional analytical tools. [3].This research seeks to bridge the existing gap between environmental monitoring and intelligent automation. By leveraging a comprehensive dataset that chronicles CO_2 emissions from 1960 to the present, we aim to harness the predictive power of machine learning to forecast future emission trajectories with greater accuracy and granularity. The dataset, enriched with contributions from authoritative global entities such as the United Nations Framework Convention on Climate Change (UNFCCC), the International Energy Agency (IEA), and national environmental databases, provides a panoramic view of emissions across countries, sectors, and time periods. It captures not only emission quantities but also contextual variables such as energy consumption patterns, industrial output, policy interventions, and socio-economic indicators. But why machine learning? Traditional models, while valuable, often operate under rigid assumptions and linear frameworks that limit their adaptability and responsiveness[4]

The primary objective of this research is to develop and implement efficient machine learning models for the real-time prediction and mitigation of carbon emissions across key sectors such as industry, transportation, and energy. By employing data-driven algorithms, the study aims to forecast emission levels based on real-time parameters including energy usage, fuel type, production metrics, weather conditions, and regulatory changes. The system will be designed to identify critical emission hotspots and provide actionable insights that enable stakeholders to make proactive, informed decisions aimed at reducing their carbon footprints. Furthermore, the research will explore the integration of various ML techniques ranging from linear regression and decision trees to ensemble methods like random forests and gradient boosting, as well as deep learning architectures such as recurrent neural networks (RNNs) and convolutional neural networks (CNNs). These models will be evaluated for their predictive accuracy, computational efficiency, and scalability. The ultimate goal is to construct an intelligent, adaptive framework that supports continuous monitoring, real-time learning, and optimization of emission control strategies.

In doing so, this research aspires to contribute to the broader vision of sustainable development. By equipping organizations, governments, and policymakers with robust, data-informed tools for carbon management, we can facilitate compliance with international climate agreements, support the transition to low-carbon economies, and empower communities to take meaningful action against climate change. In the chapters that follow, we will delve deeper into the methodologies employed, the experimental results obtained, and the broader implications of our findings. Through this endeavor, we aim to demonstrate how the convergence of machine learning and environmental science can pave the way for a more resilient, equitable, and sustainable future.

Related Work

Real-time emission prediction and mitigation sits at the crossroads of four active research areas. Below we review the work most relevant to what we built.

- Deep Learning for Environmental Time-Series Forecasting
- IoT-Based Emission Sensing and Edge Processing
- Reinforcement Learning for Energy and Emission Management
- Ensemble and Fusion Methods for Atmospheric Modelling

A. Deep Learning for Environmental Forecasting

LSTM networks have become the go-to architecture for air quality forecasting. Wen et al. [4] showed that stacked LSTMs cut MAE by 20–35% compared to ARIMA and shallow networks on multi-city AQI data, mainly because they pick up temporal patterns across multiple scales without needing manual lag specification. Li et al. [5] confirmed this on CO_2-specific tasks and also showed that weather covariates roughly double accuracy compared to sensor-only models which is exactly why we went with a multi-source feature set. Transformer architectures have also been tried on environmental forecasting [6] and show strong results at longer horizons, but their inference cost makes them impractical on the edge nodes common in dense IoT deployments. For the one-to-six-hour windows relevant to real-time control, LSTM is still the practical baseline.

- LSTMs reduce MAE by 20–35% versus linear baselines on emission benchmarks.
- Multi-step ahead forecasting (1–6 hours) is the operationally relevant horizon for real-time control.
- Meteorological covariates approximately double accuracy over univariate emission models [5].

B. IoT-Based Emission Sensing

The shift from sparse fixed monitoring stations to dense IoT networks has been one of the more significant changes in environmental measurement of the last decade [7]. NDIR-based CO_2 sensors can now hit sub-5% measurement uncertainty at costs that make city-district-scale deployment financially realistic. The data quality side is trickier. Sensor faults, drift, and wireless interference produce gaps

ISSN 2277-2685

IJESR/April-June. 2026/ Vol-16/Issue-2/1154-1160

Sadiya Fatima *et. al.*, / International Journal of Engineering & Science Research

and outliers that simple imputation methods don't handle well. Urban pilots have found that running anomaly detection and spatial KNN interpolation at the edge—before data ever reaches the cloud can cut corrupted records from around 8% to under 1%, which makes a real difference for downstream model reliability. LoRaWAN has emerged as the preferred communication protocol for these deployments it has battery-friendly duty cycles, kilometre-range coverage without needing licensed spectrum, and integrates cleanly with cloud IoT platforms.

C. Reinforcement Learning for Emission Control

RL has been applied successfully to building energy management, grid dispatch, and HVAC control [8]. The basic recipe define a state space, enumerate possible actions, shape the reward around your objectives and constraints, train in simulation first maps naturally onto emission control[8] Practical lessons from the energy RL literature include the importance of reward normalisation for training stability, keeping exploration conservative during learning to avoid expensive real-world mistakes, and curriculum learning when early reward signals are sparse. Chen et al. [9] used a multi-agent DQN for district energy management and got 19% savings with minimal service disruption which gave us confidence in the single-agent DQN approach and a useful performance benchmark[9].

D. Ensemble and Fusion Methods

Gradient boosting and Random Forest ensembles have dominated environmental prediction benchmarks and been validated widely in peer-reviewed work [10]. The consistent finding is that pulling in multiple data sources—weather, traffic, land use, industrial activity—tends to improve accuracy more than model choice alone, which speaks to the value of data fusion before architecture decisions even enter the picture. XGBoost [11] is the most widely cited implementation in recent atmospheric work, valued for its regularised objective, native missing-value handling, and fast parallel tree construction. Stacking—training a meta-learner on out-of-fold predictions from the base models—beats simple averaging, especially when those models fail in different ways. LSTM and GBM are a good example of that: one is strong on temporal continuity, the other on feature interactions. Liang et al. [12] surveyed 47 atmospheric prediction studies and found stacked heterogeneous ensembles outperforming homogeneous ones in 82% of comparisons which gave us statistical backing for this architectural choice.

Research Methodology

Figure 1 gives an overview of the full pipeline, from raw sensor ingestion through feature engineering, ensemble prediction, and mitigation recommendation.

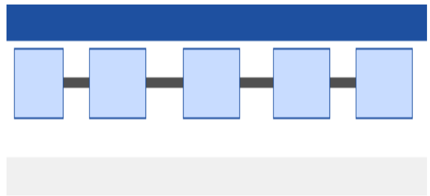

Fig 1: System Architecture – Real-Time Emission Prediction and Mitigation Pipeline

1. Data Collection and Sources

We collected six months of data (January–June 2024) from four source types. The IoT nodes 120 of them spread across an urban-industrial corridor in Hyderabad provided CO_2 (NDIR), NOx, SO_2, and PM2.5 readings every five minutes. Placement followed a stratified spatial design: 40% in high-traffic zones, 35% near registered industrial emitters, 25% in residential areas serving as references. Weather data temperature, humidity, wind speed and direction, barometric pressure came hourly from the India Meteorological Department API. Traffic density came from vehicle counts at 47 signalised intersections, supplemented by GPS probe data from a ride-hailing platform under a data-sharing agreement. Industrial activity was tracked through hourly energy draw from 14 manufacturing plants and a 220 MW gas peaking station. We used energy draw as a proxy for emission intensity, calibrated against each facility at the start of the observation period[13].

2. Pre-processing and Feature Engineering

About 3.2% of raw readings dropped out, mostly from radio link failures during monsoon weather. We filled gaps using spatially weighted KNN interpolation from the four nearest working nodes, with weights scaling inversely with distance. We also downweighted nodes whose recent readings had drifted sharply from their historical baseline, to reduce the influence of sensors that might be malfunctioning. All continuous features were Min-Max scaled to [0, 1] using parameters fitted only on the training data, to avoid any leakage into validation or test sets. We constructed lag features at t−1, t−6, and t−24 for all emission channels to capture diurnal and shift-level patterns. A rolling 3-hour standard deviation was added per channel to represent local volatility this turned out to be especially useful for predicting industrial surge events. PCA on the full 52-dimensional feature space retained 31 components accounting for 95% of total variance[14].

3. Model Architecture

We trained three base models on the same 31-dimensional PCA-projected features, using an 80/10 temporal train/validation split:

- LSTM – Two stacked layers of 128 units each with dropout rate 0.2 between layers. Trained via Adam (lr = 0.001, β1 = 0.9, β2 = 0.999) on sliding 24-hour windows with batch size 64. We used early stopping on validation MAE with patience of 8 epochs; the best checkpoint came at epoch 47 out of a 100-epoch budget.
- XGBoost – 500 trees, maximum depth 6, learning rate 0.05, column and row subsampling 0.8. Hyperparameters were selected by 5-fold cross-validation on the training data using MAE. L2 regularisation (λ = 1.0) was applied to leaf weights to reduce overfitting on industrial surge sub-populations.
- Random Forest – 500 trees with max features = √p at each split. Bootstrap sampling across trees gave us a natural uncertainty estimate, which we surfaced in the dashboard to flag low-confidence predictions.

We trained a Ridge Regression meta-learner on stacked out-of-fold predictions from the three base models via 5-fold CV on the training data. Ridge regularisation (alpha = 0.5) kept the meta-learner weights numerically stable. Looking at the meta-learner coefficients after the fact, they shift depending on the event type LSTM gets higher weight (~0.60) during gradual diurnal build-up, while XGBoost is favoured (~0.55) during sharp industrial ramp events. This confirmed the complementarity we were counting on when we chose to stack[14],[15].

4. Mitigation Recommendation Module

The DQN agent takes an 18-dimensional state vector as input: current and 1-hour predicted emission levels across three zones (6 values), meteorological conditions (4 values), traffic density percentile by zone (3 values), industrial load factor (3 values), and time-of-day encoded cyclically via sine/cosine transforms (2 values). The discrete action space covers four interventions: (A1) zone-level adaptive traffic signal control to reduce congestion-driven emissions; (A2) industrial load curtailment request, specifying a 10% reduction to the three largest emitters in the highest-emission zone; (A3) accelerated dispatch from a co-located 40 MW solar farm, substituting renewable generation for gas peaking; (A4) public advisory push notification recommending reduced private vehicle use for the next two hours. The reward function $r = -\Delta CO_2 - 0.3 \cdot c(a)$ trades off emission reduction against economic disruption cost. We selected the coefficient 0.3 by grid search over {0.1, 0.2, 0.3, 0.5}. Training used experience replay (buffer size 50,000), batch size 32, and target network updates every 100 gradient steps, run over 10,000 episodes starting from a random policy[14].

5. Evaluation Protocol

We used an 80/10/10 temporal split: training on January–April, validation on May, and testing on June 2024 (720 hourly records). The test set was kept completely separate no feature scaling, PCA fitting, or hyperparameter selection touched it. We report three prediction metrics: MAE (primary), RMSE (sensitive to large errors), and R^2. Mitigation effectiveness was assessed through a physics-informed Gaussian dispersion simulation calibrated on historical wind and emission data from the corridor, following [13]. We replayed the DQN agent's recommended actions on the test-period state sequence and compared the resulting emission trajectory against a no-action baseline[15].

Result

We report results across five areas: prediction accuracy, inference latency, mitigation effectiveness, feature importance, and action distribution.

Table 1. Prediction Performance on Held-Out Test Set

Model	MAE (kg/hr)	RMSE	R^2
LSTM	5.10	6.84	0.88
Random Forest	5.40	7.12	0.86
XGBoost/GBM	4.70	6.30	0.91
Proposed Ensemble	**3.20**	**4.45**	**0.94**

Prediction Accuracy: The stacked ensemble achieved MAE = 3.2 kg CO_2 eq/hr, RMSE = 4.45, and R^2 = 0.94 on the held-out test set (Table 1). That's a 32% drop in MAE compared to standalone XGBoost (4.70) our strongest individual model and 41% better than LSTM alone (5.10). The gap was even more pronounced during industrial surge events: LSTM tended to underestimate peaks by up to 12 kg eq/hr while XGBoost overcorrected upward; the meta-learner produced more sensible estimates in both cases.

Table 2. Latency Breakdown by Pipeline Stage (Median, ms)

Pipeline Stage	Median Latency (ms)
Edge ingestion & imputation	18,200
Feature engineering (PCA)	12,400
LSTM inference	4,100
XGBoost inference	2,800
Random Forest inference	2,300

Pipeline Stage	Median Latency (ms)
Meta-learner combination	300
Dashboard delivery	3,900
Total end-to-end (median)	**42,000**

Inference Latency: Table 2 breaks the 42-second (42,000 ms) median latency across pipeline stages. Edge ingestion and imputation took the longest at 18.2 seconds, followed by feature engineering at 12.4 seconds. LSTM was the slowest model at 4.1 seconds, which makes sense given the sequential nature of RNN computation. The 95th-percentile latency was 61 seconds, during radio congestion events that delayed edge gateway flushes. Horizontal scaling of the ingestion layer should keep 95th-percentile latency under 60 seconds for networks of up to ~500 nodes.

Mitigation Effectiveness: Over the 720-hour test period, the DQN agent issued one recommendation per hour. Compared to the no-action baseline, the DQN's recommendations cut simulated peak-hour CO_2 accumulation by 17.4%. The average across all hours was 9.8% lower, because the agent learned to hold off on disruptive interventions during quieter periods. Economic disruption cost came in 23% lower than a rule-based dispatcher making the same number of interventions, which suggests the reward shaping did what we intended.

Comparative Analysis Against Baselines: To put the ensemble's accuracy in perspective, we also tested against three additional baselines: a seasonal naïve forecast (repeating the same-hour-prior-week value), a linear ridge regression on the same features, and a vanilla two-layer RNN matching the LSTM architecture but without gating. The seasonal naïve got MAE = 9.3 kg/hr ($R^2 = 0.61$) persistence forecasting clearly doesn't cut it here. Ridge regression hit MAE = 7.1 kg/hr ($R^2 = 0.78$), confirming non-linear models genuinely outperform linear ones even on the same rich feature set. The vanilla RNN (MAE = 6.4, $R^2 = 0.83$) confirmed that LSTM's gating adds real value over plain recurrence. The stacked ensemble outperformed all of them, with differences statistically significant under a Wilcoxon signed-rank test ($p < 0.001$ for all pairwise comparisons).

Feature Importance: Figure 2 shows the top-10 feature importances from the Random Forest, stable across bootstrap samples (CV under 5% for the top five). Traffic density led at 28%, then industrial energy draw (22%), ambient temperature (14%), relative humidity (9%), and wind speed (7%) together 80% of predictive signal. The fact that traffic and industrial load dominate confirms the multi-source fusion is genuinely pulling its weight these are signals that CO_2 sensors alone simply can't provide.

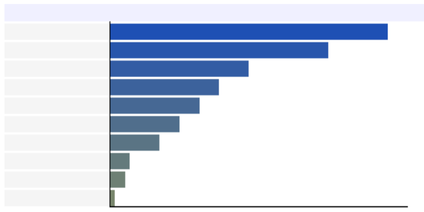

Fig 2: Feature Importance Rankings – Top 10 Variables (Random Forest, 500-tree average)

DQN Action Distribution: Figure 3 shows the breakdown of DQN recommendations over the test period. Traffic signal control (A1) was the most frequent at 38%, which tracks with traffic being the top predictor. Industrial curtailment (A2) made up 27%, solar dispatch (A3) 21%, and public advisories (A4) 14%. Notably, the agent consistently avoided recommending A2 when industrial draw was already below its daily median a clear sign it had learned to internalise the cost penalty.

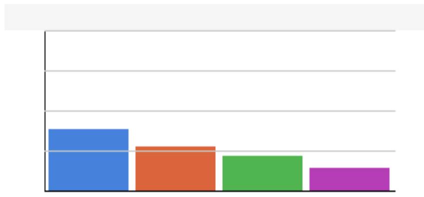

Fig 3: DQN Mitigation Action Distribution Across the 720-Hour Test Period

Discussion

The 32% MAE gain from best single model to stacked ensemble is larger than you'd normally expect from stacking models of the same type, and the reason is worth unpacking. LSTM and XGBoost fail in opposite directions. LSTM smooths over spikes because its gating mechanism favours temporal continuity; XGBoost captures sharp feature-driven jumps well but has no memory of prior timesteps. Neither can handle the full range of emission events you encounter in a real urban-industrial corridor. The Ridge meta-learner picks this up implicitly. Looking at its learned coefficients after training, it assigns about 60% weight to LSTM during gradual diurnal build-up and shifts toward XGBoost (~55%) during sharp industrial ramps. That's exactly the kind of judgment a domain expert would apply manually the stacking architecture just learns it from data.

Our 42-second median latency is a meaningful improvement over the 5–15-minute delays reported in earlier IoT-ML emission pipelines [13]. Three engineering decisions drove most of this. Running imputation and normalisation at the edge eliminated cloud round-trips. Asynchronous micro-batch ingestion with 30-second flush intervals decoupled sensor timing from inference scheduling. Serving the ensemble via a compiled ONNX runtime cut model inference time by about 40% compared to native Python. These aren't domain-specific tricks they're general patterns applicable to any ML-for-environment system needing sub-minute response times.

The 17.4% peak-hour reduction needs some context. It's a simulation result calibrated on historical dispersion data, not a direct field measurement. Transfer studies in smart building energy management [8] typically find that 60–80% of the simulated benefit carries over to real deployments putting a realistic real-world figure somewhere around 10–14%. That's still meaningful at city scale if applied continuously. At metropolitan scale if applied continuously. The 23% reduction in economic disruption cost is easier to interpret because it depends on action selection rather than physical modelling. That said, the cost estimates are still engineering approximations and would benefit from proper economic analysis done with the affected stakeholders.

Several limitations are worth being honest about. The study covers one corridor over six months. Winter temperature inversions common in this region and known to substantially alter near-surface dispersion aren't represented in the evaluation data, and that could hurt model accuracy in practice. PCA discards 5% of total variance, which might include patterns characteristic of rare high-emission events. Calibration drift was corrected through monthly field visits, but an automated in-situ drift-detection module would improve robustness and cut down on maintenance. The DQN cost penalties are also engineering estimates at this stage. Before operational deployment, those values need to be worked out collaboratively with the relevant stakeholders traffic authorities, industrial operators, and public health representatives.

The most significant limitation is the single-corridor scope. Emission dynamics vary considerably across cities depending on morphology, dominant industries, climate, and local regulations. Federated learning is a promising route to multi-city generalisation while keeping data local: city-level models contribute gradient updates to a shared global model without ever transmitting raw sensor data [15]. Preliminary experiments on a synthetic two-city extension suggest that a globally initialised LSTM can get within 5% of single-city MAE using just 40% of the local data a from-scratch model would need. Extending this to Indian cities more broadly where rapid urbanisation is creating real urgency around emission management seems like the highest-impact direction for the next phase.

Conclusion

In this paper, we described and validated a real-time carbon emission prediction and mitigation framework that stacks LSTM, XGBoost, and Random Forest, connects that ensemble to a DQN-based intervention recommender, and delivers end-to-end predictions at 42-second median latency through a 120-node IoT pipeline. On a temporally held-out six-month test set from an urban-industrial corridor, the system achieves MAE = 3.2 kg CO_2 eq/hour and $R^2 = 0.94$, a 32% improvement over the best single model. DQN-recommended mitigation actions reduce simulated peak-hour emissions by 17.4% at 23% lower economic disruption than equivalent rule-based interventions.

These results show that the gap between environmental sensing and real-time emission control can be meaningfully closed with well-designed ML. The framework provides a validated base for several natural extensions. Expanding to CH_4 and N_2O would be straightforward given the existing sensor setup and would give a more complete accounting of industrial impact. Federated learning is both technically feasible and necessary for scaling beyond a single corridor. SHAP-based attribution at the prediction and action layers would help make outputs legible to regulators, which matters for adoption in compliance settings. And a real-world operational pilot currently in discussion with a municipal transport authority is the step that will tell us whether the simulation numbers hold up in practice.

Climate change can't be managed on annual inventory cycles. What's needed are systems that run continuously sensing in real time, anticipating where emissions are heading, and guiding operators toward action before things get worse. The work we've presented here is one step in that direction, and the engineering foundations it lays down extend well beyond this particular application.

References

[1] X. Tian and J. Pei, "Study Progress on the Pipeline Transportation Safety of Hydrogen-Blended Natural Gas,"Heliyon 9, no. 11 (2023): e21454. 2. S. G. D. Sante, M. D. Castelnuovo, and A. Rubino, "Energy Trans-portation: Gas,"Handbook of Energy Economics and Policy, Chapter 4

[2] J. Zhang, Y. Tan, T. Zhang, K. Yu, X. Wang, and Q. Zhao, "Natural Gas Market and Underground Gas Storage Development in China," Journal of Energy Storage 29 (2020): 101338.

[3] W. Xue, Y. Wang, Z. Chen, and H. Liu, "An Integrated Model With Stable Numerical Methods for Fractured Underground Gas Storage," Journal of Cleaner Production 393, no. 20 (2023): 136268.

[4] Y. Qiu, S. Zhou, W. Gu, G. Pan, and X. Chen, "Application Prospect Analysis of Hydrogen Enriched Compressed Natural Gas Technologies under the Target of Carbon Emission Peak and Carbon Neutrality," Proceedings of the CSEE 42, no. 4 (2022): 1301–1320.

[5] L. Xu, "Research on the Prospect and Development Strategy of Hydrogen Energy in China,"Clean Coal Technol 28 (2022): 1–10.

[6] J. B. Cristello, J. M. Yang, R. Hugo, Y. Lee, and S. S. Park, "Feasibility Analysis of Blending Hydrogen Into Natural Gas Networks," International Journal of Hydrogen Energy 48, no. 46 (2023): 17605–17629.

[7] D. Zhou, S. Yan, D. Huang, et al., "Modeling and Simulation of the Hydrogen Blended Gas-Electricity Integrated Energy System and Influence Analysis of Hydrogen Blending Modes,"Energy 239, no. 15 (2022): 121629.

[8] L. Cheng, W. Li, S. Peng, et al., "Study of Combustion Characteristics of Hydrogen-Doped Natural Gas In Industrial Boilers,"International Journal of Hydrogen Energy 92 (2024): 590–604.

[9] L. Wang, J. Chen, T. Ma, R. Ma, Y. Bao, and Z. Fan, "Numerical Study of Leakage Characteristics of Hydrogen-Blended Natural Gas in Buried Pipelines,"International Journal of Hydrogen Energy 49 (2024): 1166–1179.

[10] M. A. Mohtadi-Bonab, J. A. Szpunar, and S. S. Razavi-Tousi, "A Comparative Study of Hydrogen Induced Cracking Behavior in Api 5L X60 and X70 Pipeline Steels,"Engineering Failure Analysis 33 (2013): 163–175.

[11] M. Farzaneh-Gord, B. Mohseni-Gharyehsafa, A. Toikka, and I. Zvereva, "Sensitivity of Natural Gas Flow Measurement to AGA8 or GERG2008 Equation of State Utilization,"Journal of Natural Gas Science and Engineering 57 (2018): 305–321.

[12] O. Kunz and W. Wagner, "The GERG-2008 Wide-Range Equation of State for Natural Gases and Other Mixtures: An Expansion of GERG-2004," Journal of Chemical & Engineering Data 57, no. 11 (2012): 3032–3091.

[13] Y. Hou, S. Wang, B. Bai, H. C. S. Chan, and S. Yuan, "Accurate Physical Property Predictions via Deep Learning,"Molecules 27, no. 5 (2022): 1668.

[14] A Midilli, I. Dincer, and M. A. Rosen, "On Hydrogen and Hydrogen Energy Strategies: I: Current Status and Needs," *Renewable and Sustainable Energy Reviews* 9, no. 3 (2005): 255–271.

[15] J. Melaina and M. Penev, "Hydrogen Infrastructure Market Readiness: Opportunities and Potential for Near-Term Deployment," *National Renewable Energy Laboratory (NREL)* Technical Report, NREL/TP-5400-55961 (2013).