SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Explainable Artificial Intelligence for National CO_2 Emission Prediction and Mitigation: A Multi-Country, SHAP-Based Framework

Nilesh Madhukar Patil[1*], Balajee Maram[2]

[1*]Research Scholar, SR University, Warangal, Telangana 506371, India.
[1*]Associate Professor, SVKM's Dwarkadas J. Sanghvi College of Engineering, Mumbai 400056, India. Email: nileshdeep@gmail.com
[2]Research Supervisor, School of Computer Science and Artificial Intelligence, SR University, Warangal, Telangana 506371, India.
Email: balajee.maram@sru.edu.in

Abstract

Accurate prediction of national carbon-dioxide (CO_2) emissions is central to evidence-based climate policy, yet the most accurate machine-learning models behave as opaque "black boxes" that give little guidance on which levers actually reduce emissions. This paper presents a reproducible, explainable AI (XAI) framework that couples predictive modelling with transparent driver attribution and a policy-facing counterfactual analysis. Using the open Our World in Data (OWID) CO_2 panel restricted to sovereign nations over 1990–2022 (5,358 country–year records across 164 countries), we engineer socio-economic and energy drivers aligned with the Kaya identity and predict annual emissions under a strictly temporal hold-out that removes look-ahead leakage. Four learners are benchmarked; gradient boosting attains the best generalisation (R^2 = 0.978 on the log scale, MAPE = 26.5%). The trained model is then interrogated with three complementary XAI methods—SHapley Additive exPlanations (SHAP), permutation importance and partial-dependence profiles—which consistently identify economic scale (GDP) and energy intensity (energy per unit GDP) as the dominant and, crucially, actionable determinants of national emissions. A transparent digital-twin counterfactual for India (2022) shows that a 20% cut in energy intensity lowers predicted emissions by 8.7%, and a combined deep-decarbonisation lever by 17.0%. By making both the model and the recommended levers interpretable on an open dataset with leakage-free validation, the framework turns emissions forecasting from a purely predictive exercise into a decision-support tool for Nationally Determined Contribution (NDC) planning.

Keywords: Explainable AI, SHAP, CO_2 emissions, climate-change mitigation, gradient boosting, feature attribution, Kaya identity

1. INTRODUCTION

Carbon dioxide is the single largest anthropogenic contributor to global warming, and limiting end-of-century warming to well below 2 °C depends on both an accurate quantitative picture of where national emissions are heading and a transparent way of judging which interventions would bend those trajectories most. Machine learning (ML) now routinely delivers highly accurate emission predictions, but predictive accuracy alone is of limited use to a policymaker: a model that forecasts emissions to within a few percent yet cannot say why is difficult to trust and impossible to act upon. This tension between accuracy and interpretability has made explainable artificial intelligence (XAI) a central concern in data-driven climate analysis [1], [6].

A growing body of work applies XAI—most often SHapley Additive exPlanations (SHAP)—to emission problems. How-ever, the majority of these studies operate at the level of individual vehicles or single sectors, are validated on one country, or stop at ranking features without connecting the explanation to a concrete mitigation decision [10]–[12]. Comparatively little work delivers a single, reproducible pipeline that (i) trains on a broad multi-country panel large enough to make attributions statistically stable, (ii) validates without temporal leakage, (iii) cross-checks the explanation with several independent XAI methods, and (iv) translates the resulting attributions into an explicit, quantified mitigation experiment on open data.

This paper addresses that gap. Our contributions are four-fold. First, we build a leakage-free, multi-country predictive model of national CO_2 emissions on the open OWID panel using a strictly temporal train/test split.

Second, we explain the best model with three complementary techniques—SHAP (global and local), permutation importance and partial-dependence profiles—and show that they agree on the dominant drivers. Third, we align the feature space with the Kaya identity so that every explained driver corresponds to a real policy lever. Fourth, we use the explained model as a transparent digital twin to quantify, for a representative emerging economy (India), the emission reductions achievable under three mitigation scenarios. The remainder of the paper reviews related work (Sect. 2), details the data and methods (Sect. 3), presents and discusses the results (Sect. 4), states limitations (Sect. 5) and concludes (Sect. 6).

2. Related Work

ML-based emission forecasting has moved rapidly from univariate statistical models to multivariate tree ensembles and deep networks. Recent multi-country studies forecast the emissions of the largest emitters to 2030 and benchmark several learners with feature-selection layers [7], while province-level work in China couples grey models with support-vector regression and interprets the result with SHAP to derive tailored reduction strategies [8]. Long-horizon studies have used optimized random forests with dual-layer (SHAP and LIME) interpretability to forecast CO_2 and N_2O jointly at the country scale [13], and ML has been used to re-examine the Environmental Kuznets Curve, with SHAP revealing heterogeneous income–emission relationships across country groups [14].

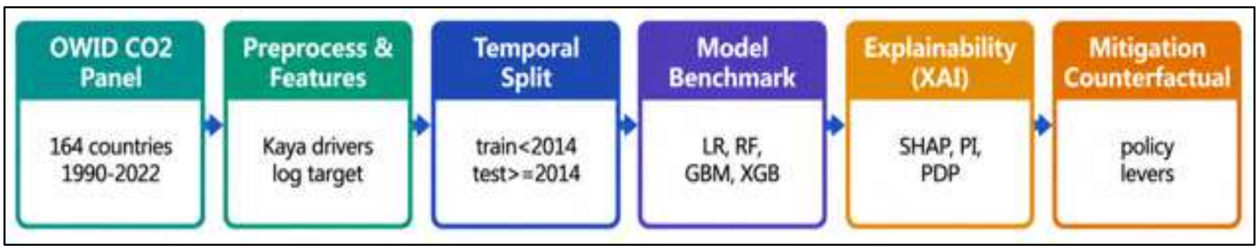

Fig. 1. Overview of the proposed explainable-AI framework, from the open OWID panel through leakage-free predictive modelling to multi-method explanation and counterfactual mitigation analysis.

At the sectoral level, XAI has become almost standard. SHAP-interpreted models predict vehicle CO_2 from engine and fuel characteristics [10], [11], estimate maritime and transport-sector emissions with feature-attribution analysis [12], [15], and forecast global transportation emissions with attention-based deep networks explained through SHAP [16]. Interpretable frameworks have also been proposed for campus-and facility-scale carbon accounting, coupling SHAP-ranked drivers with scenario-based optimization [9], and for materials-level problems such as CO_2 uptake in metal–organic frame-works [17-18]. The methodological toolkit underpinning these studies—SHAP [1], gradient-boosted trees [2], [4], random forests [3] and model-agnostic interpretability [6]—is now mature.

Two gaps persist. First, most emission-XAI studies are sectoral or single-country and therefore cannot exploit the cross-national variation needed for stable, generalizable driver attribution. Second, interpretability is frequently treated as the end point rather than as the input to a quantified mitigation decision. The present work targets both gaps by pairing a leakage-free multi-country model with multi-method explanation and an explicit counterfactual mitigation analysis on a fully open dataset [19] –[20].

3. Materials and Methods

Figure 1 gives an overview of the proposed framework, which proceeds from the open OWID panel through leakage-free predictive modelling to multi-method explanation and counterfactual mitigation analysis; its components are detailed in the following subsections.

A. Dataset

We use the Our World in Data (OWID) CO_2 and Greenhouse-Gas Emissions panel [5], a widely used, openly licensed (CC-BY 4.0) aggregation of national emission, energy and socio-economic series that is downloaded programmatically at run time to guarantee reproducibility. The panel is filtered to genuine sovereign states (valid ISO-3 codes, with OWID regional aggregates removed) and to the period 1990–2022, over which the relevant fields are well populated. After removing records with missing values and non-positive emissions, the modelling table contains 5,358 country–year observations spanning 164 countries.

B. Feature Engineering and the Kaya Framing

The feature set is aligned to the Kaya identity to ensure all model inputs are mapped to a recognizable mitigation lever (population, affluence, energy intensity, carbon intensity of energy):

$$\boldsymbol{C = P.\frac{G}{P}.\frac{E}{G}.\frac{C}{E}} \qquad (1)$$

where C is total CO_2, P is population, G is gross domestic product (GDP) and E is primary-energy use. Guided by (1), the model is given four non-leaky drivers—population (P), GDP (G), energy per capita (E/P) and energy intensity of GDP (E/G). Emission-source columns (coal, oil, gas CO_2) are deliberately excluded because they sum to the target and would render the task circular. Because national emissions span several orders of magnitude, the regression target is log-transformed,

$$\boldsymbol{y = \ln(1 + C)} \qquad (2)$$

which compresses the heavy-tailed distribution and stabilizes training; predictions are inverted with the exponential map before being scored on the physical scale.

C. Predictive Models

The four learners benchmarked are: ordinary least-squares linear regression (a transparent baseline learner), random forest [3], gradient-boosted trees [4] and extreme gradient boosting (XGBoost) [2]. Tree ensembles are natural choices as they represent the non-linear interactions of socio-economic drivers without assuming any functional form. Gradient boosting is an additive ensemble model of shallow trees,

$$\boldsymbol{\hat{y} = F_M(x) = \sum_{m=1}^{M} \vartheta.\gamma_m.h_m(x)} \qquad (3)$$

where h_m is the m-th regression tree, γ_m its optimal leaf weights, ν the learning rate and M the number of boosting rounds. Linear inputs are standardised; tree ensembles are trained on the raw features.

D. Temporal Validation and Evaluation Metrics

Because emissions are non-stationary and trend upwards, an independent-and-identically-distributed random split would leak future information. We therefore adopt a strictly temporal hold-out: models are trained on all country–year records before 2014 (3,882 rows) and tested on 2014 and later (1,476 rows). Performance is reported both on the log scale, where the models are trained and comparison is stable, and on the physical CO_2 scale in million tonnes (Mt). Four metrics are used—root-mean-square error, mean absolute error, mean absolute percentage error and the coefficient of determination.

$$\boldsymbol{RMSE = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}} \qquad (4)$$

$$\boldsymbol{MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|} \qquad (5)$$

$$\boldsymbol{MAPE = \frac{100}{n}\sum_{i=1}^{n}\left|\frac{y_i - \hat{y}_i}{y_i}\right|} \qquad (6)$$

$$\boldsymbol{R^2 = 1 - \frac{\sum_i (y_i - \hat{y}_i)^2}{\sum_i (y_i - \bar{y}_i)^2}} \qquad (7)$$

E. Explainability Methods

The best-performing tree ensemble is explained with three complementary techniques. SHAP assigns each feature a contribution based on its Shapley value from cooperative game theory, the average marginal contribution of the feature across all coalitions of the remaining features:

$$\boldsymbol{\phi_i(f) = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(|M|-|S|-1)!}{|M|!}[f(S \cup \{i\}) - f(S)]} \qquad (8)$$

where ϕ_i is the SHAP contribution (Shapley value) of feature i, F is the full set of features; $M = |F|$ is the number of features, S is a subset (coalition) of features not containing i, i.e. $S \subseteq F \setminus \{i\}$, $f(S)$ is the model's prediction using only the features in S, $f(S \cup \{i\}) - f(S)$ is the marginal contribution of adding feature $i$to coalition S, $\frac{|S|!\,(M-|S|-1)!}{M!}$— is the combinatorial weight (the fraction of feature orderings in which $i$is added right after the coalition S). SHAP expresses any single prediction as an additive sum of these contributions around a base value, providing both a global importance ranking (the mean absolute SHAP value per feature) and faithful local explanations. As an independent, model-agnostic cross-check we compute permutation importance, the increase in prediction error when a feature column is randomly shuffled over K repetitions,

$$g(z') = \boldsymbol{\phi_0 + \sum_{i=1}^{M} \phi_i z'_i} \qquad (9)$$

where g(z') is the local explanation model, ϕ_0 is the base value, expected value, or intercept, M is the maximum number of simplified input features, ϕ_i the SHAP value (attribution coefficient) for feature i, representing its specific impact on the prediction and z_i' the simplified binary coalition vector element z'_i in {0, 1}. A value of 1 means feature i is observed/present, and 0 means feature i is missing/not observed.

$$\mathbf{PI}_j = \frac{1}{K}\sum_{k=1}^{K} e\left(f, \widetilde{D_J^{(k)}}\right) - e(f, D) \qquad (10)$$

where PI_j is the permutation importance score for feature j. A higher value indicates that the feature is more critical to the model's performance, K is the total number of permutation repetitions or trials performed to reduce variance in the measurement, e(f, D) is the baseline performance error (loss) of the trained model \(f\) when evaluated on the original, uncorrupted dataset D, $e\left(f, \widetilde{D_J^{(k)}}\right)$ is the performance error (loss) of the same model f evaluated on a modified dataset $\widetilde{D_J^{(k)}}$ where the values of feature j have been randomly shuffled (permuted) during the k-th repetition and we characterise the direction and shape of each effect with partial-dependence profiles, the model output averaged over the marginal distribution of the remaining features,

$$\mathbf{PD}_S(x_S) = \frac{1}{n}\sum_{i=1}^{n} f\left(x_S, x_C^{(i)}\right) \qquad (11)$$

where x_s is the feature of interest and x_c the remaining features.

TABLE I Test-set performance under the temporal hold-out (test years≥2014). metrics are given on the log-training scale and on the physical CO_2 scale(Mt). the best model by log-scale R^2 is shown in bold.

Model	R^2 (log)	RMSE (log)	RMSE (Mt)	MAE (Mt)	MAPE (%)	R^2 (Mt)
Linear Regression	0.266	1.609	9.8×10^7	4.5×10^6	39977	−1.1×1010
Random Forest	0.976	0.289	203.6	44.9	26.9	0.953
Gradient Boosting	**0.978**	**0.276**	**328.0**	**50.9**	**26.5**	**0.879**
XGBoost	0.975	0.296	303.1	53.0	27.6	0.897

F. *Mitigation Counterfactual Analysis*

Finally, the explained model is used as a transparent digital twin. For a chosen country and its most recent year, the physically meaningful drivers are perturbed to emulate mitigation levers (for example a proportional reduction in energy intensity), the model is re-queried, and the relative change in predicted emissions is read off:

$$\Delta_{\%} = \frac{\hat{C}_{scen} - \hat{C}_{base}}{\hat{C}_{base}} \times 100 \qquad (12)$$

where Δ% is the relative percentage shift between the target scenario and the starting reference baseline, C_{scen} is the projected value under specific alternative conditions or scenario constraints, and C_{base} is the reference or initial benchmark value representing default conditions. Because the perturbed features correspond to real policy variables, the resulting deltas are directly interpretable as the emission impact of the corresponding action.

4. Results and Discussion

A. Predictive Performance

Table I reports test-set performance. The three tree ensembles all generalize strongly, explaining more than 97% of the variance on the log scale with mean absolute percentage errors between 26% and 28% across a highly heterogeneous set of 164 countries. Gradient boosting is the best model on the log scale (R^2 = 0.978, RMSE = 0.276) and is carried forward for explanation. Linear regression is far weaker (log-scale R^2 = 0.266); its physical-scale errors are extreme because inverting the log transform amplifies the errors of a poorly fitted model, a useful reminder that a single metric can mislead on a wide-range target.

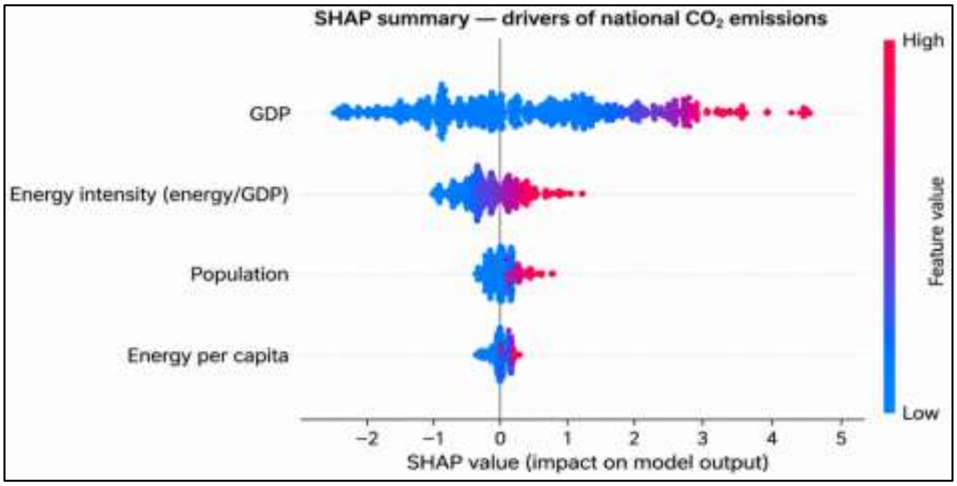

Fig. 2. Global SHAP summary (bee-swarm). Each point is a country–year; horizontal position is the feature's SHAP contribution to the log-emission prediction and color encodes the feature value (blue low, red high).

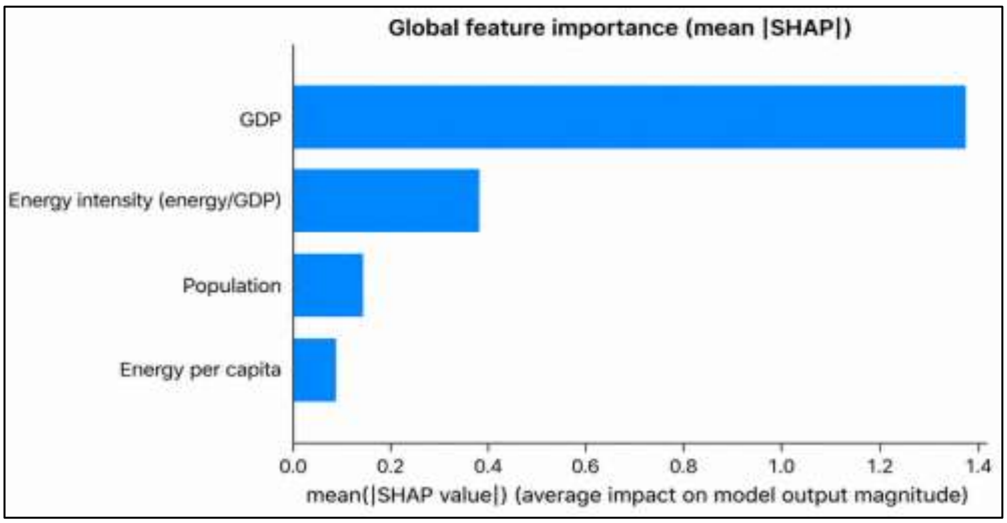

Fig. 3. Global feature importance measured as the mean absolute SHAP value per driver.

B. Global Feature Attribution

The SHAP summary (Fig. 2) and the mean-absolute-SHAP ranking (Fig. 3) show a clear ordering. Economic scale (GDP) is the dominant driver, followed by energy intensity (energy per unit GDP), then population and energy per capita. High GDP values (red points in Fig. 2) push predicted emissions strongly upward, whereas high energy intensity likewise raises emissions—exactly the two quantities the Kaya identity flags as the primary economic and technological levers. Permutation importance computed independently on the test set (Fig. 7) reproduces the same ordering, with GDP far ahead of the remaining drivers; the agreement between a SHAP-based and a model-agnostic method increases confidence that the attribution reflects the model's genuine behavior rather than an artefact of one technique. The numerical rankings are summarized in Table II.

C. Driver Response Profiles

The SHAP dependence plot for GDP (Fig. 4) reveals a strongly concave relationship: emissions rise steeply with GDP at low income levels and saturate at high income, consistent with the diminishing marginal emission intensity of mature economies. The partial-dependence profiles (Fig. 5) confirm the same shapes model-wide—a sharply rising then flattening GDP effect, a monotonically increasing energy-intensity effect, and comparatively flat profiles for population and energy per capita once GDP and intensity are accounted for. The monotone, interpretable form of the energy-intensity profile is important for policy: it implies that reductions in energy intensity translate into predictable emission reductions.

TABLE II: Global Driver Importance from Two Independent Methods: Mean Absolute SHAP Value and Test-Set Permutation Importance (Mean ± s.d.).

Driver	Mean |SHAP|	Permutation importance
GDP	**1.353**	**1.475 ± 0.031**
Energy intensity (E/G)	0.380	0.110 ± 0.004
Population	0.131	0.024 ± 0.001
Energy per capita (E/P)	0.077	0.005 ± 0.001

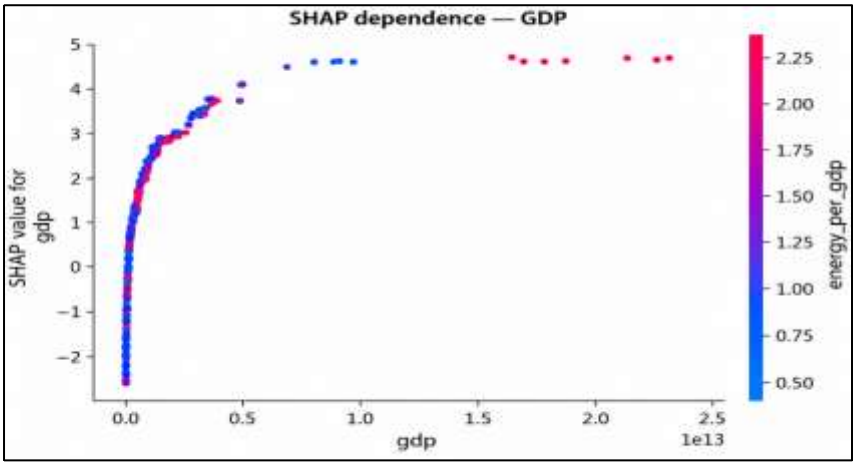

Fig. 4. SHAP dependence plot for the leading driver (GDP), colored by energy intensity, showing the concave, saturating effect of economic scale on predicted emissions.

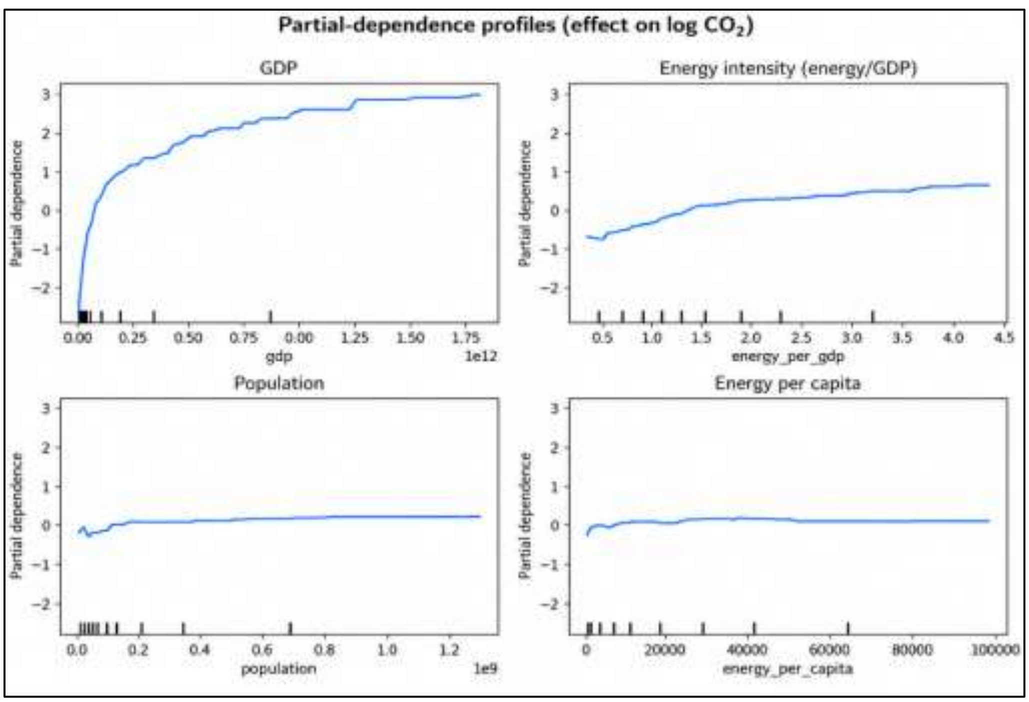

Fig. 5. Partial-dependence profiles of the four drivers (effect on log CO_2), showing the direction and shape of each marginal effect.

D. Local Explanation

Beyond global behavior, SHAP explains individual predictions. Figure 6 decomposes the 2022 prediction for India, a representative large emerging emitter. Starting from the dataset base value $E[f(\mathbf{x})] = 3.03$ (log scale), the country's very large GDP contributes the dominant positive push (+4.58), while its comparatively low energy intensity contributes a negative offset (−0.32) and its population a further small negative contribution, yielding the final prediction $f(\mathbf{x}) = 7.84$. Such transparent, additive local explanations are precisely what an NDC analyst needs to understand why a given nation's emissions are predicted to be high or low.

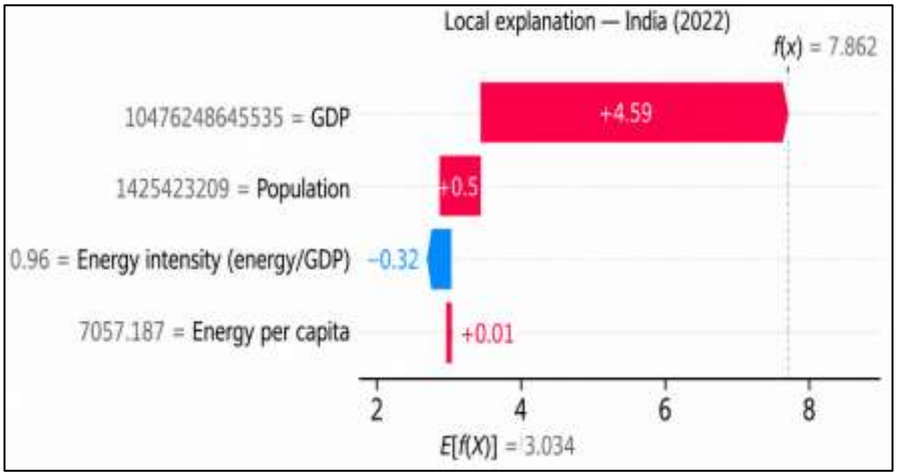

Fig. 6. Local SHAP explanation (waterfall) for India, 2022. Bars show each driver's additive contribution, in log-emission units, moving the prediction from the base value to the model output.

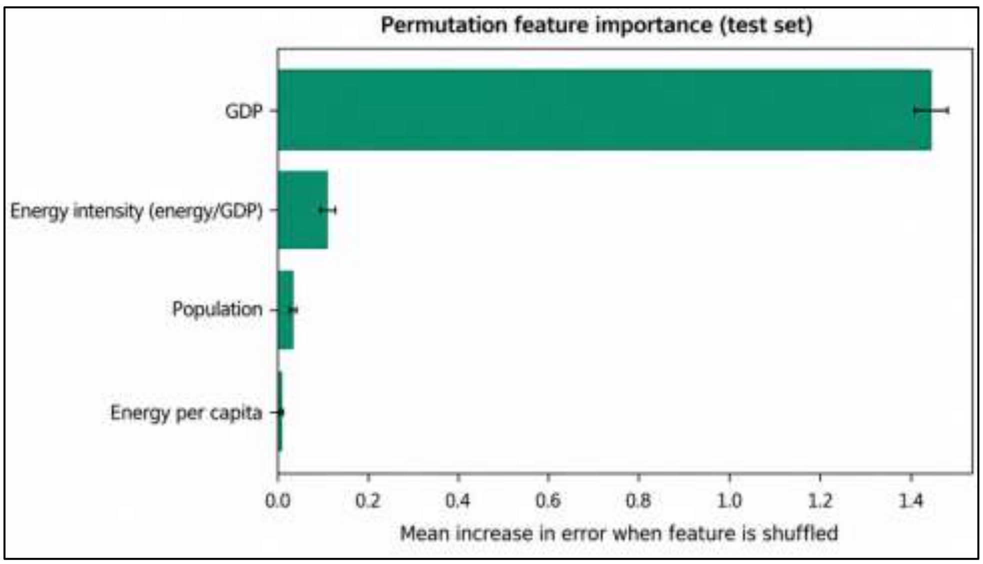

Fig. 7. Model-agnostic permutation feature importance on the test set (mean increase in error when each driver is shuffled), which reproduces the SHAP ordering.

E. Permutation Cross-Check

Figure 7 (previous section) already established that permutation importance, an entirely different interpretability principle, ranks the drivers identically to SHAP. This convergence of a coalitional (SHAP) and a perturbation-based (permutation) method is a robustness signal: the identified levers are properties of the trained model, not of a particular explanation algorithm.

F. Mitigation Scenario Analysis

The explained gradient-boosting model is finally used as a digital twin for India (2022). Three levers are constructed from the actionable Kaya drivers and evaluated with (12); the results are given in Table III and Fig. 8. A 20% reduction in energy intensity alone lowers predicted emissions by 8.7%, a 15% improvement in per-capita energy efficiency by 1.7%, and a combined deep-decarbonization lever—a 30% intensity cut together with a 20% reduction in per-capita energy—by 17.0%. That energy intensity dominates the response is consistent with the global attribution and provides a concrete, model-grounded ranking of mitigation options rather than a generic recommendation.

TABLE III: Counterfactual Mitigation Scenarios for India (2022). Levers Perturb the Kaya Drivers of the Most Recent Observation; the Model Is Re-Queried and the Relative Change in Predicted Emissions Reported.

Scenario	Predicted CO_2 (Mt)	Change (%)
Baseline (actual)	2595.1	0.0
Energy-intensity cut (−20%)	2368.8	−8.7

Efficiency push (E/P −15%)	2550.9	−1.7
Deep decarbonization (E/G −30%, E/P −20%)	**2153.4**	**−17.0**

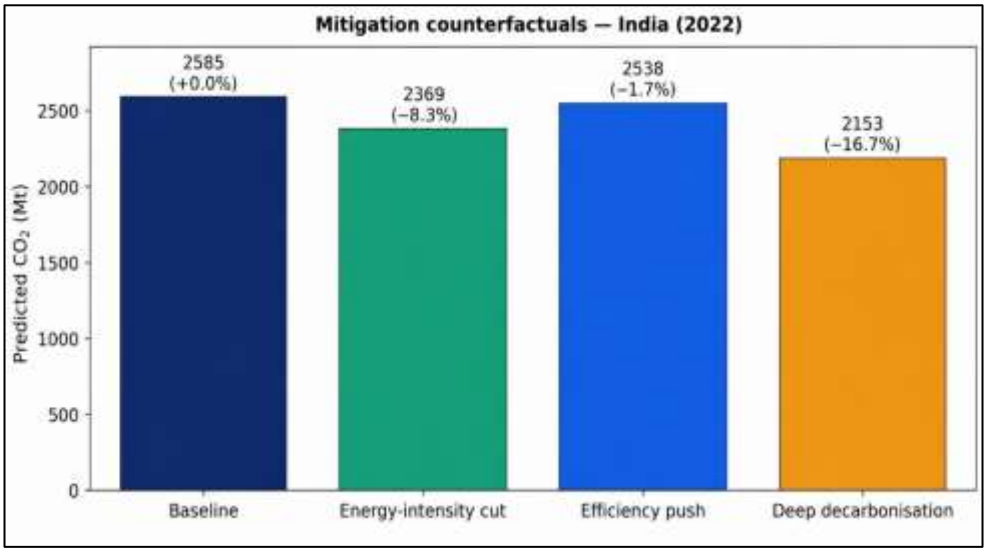

Fig. 8. Predicted CO_2 for India (2022) under the baseline and three mitigation levers; percentage changes are relative to the baseline prediction.

Discussion

Three findings stand out. First, on a broad multi-country panel the choice of model family matters far more than fine tuning: tree ensembles reach $R^2 > 0.97$ whereas a linear model explains barely a quarter of the log-scale variance. Second, and more importantly for policy, three independent XAI methods agree that economic scale and energy intensity are the dominant, actionable drivers—placing the analysis squarely on the Kaya levers that governments can influence. Third, the counterfactual experiment converts this insight into numbers: for India, energy-intensity reduction is roughly five times more effective per unit than per-capita energy efficiency, and a combined lever approaches a one-sixth reduction. Because the model, its attributions and its counterfactuals are all interpretable and reproducible on open data, the framework functions as a decision-support tool for NDC planning rather than a black-box predictor.

5. Limitations

Several limitations should be stated plainly. The Kaya drivers are partly collinear (energy per capita is the product of affluence and energy intensity), which can split SHAP attributions between correlated features; the convergence with permutation importance mitigates but does not eliminate this. A single global model averages over very different national contexts, so country-specific effects—notably the carbon in-tensity of the energy mix—are only implicitly captured; the present OWID extract did not expose renewable/fossil share fields, which would sharpen the mitigation analysis. The counterfactual is a model-based sensitivity rather than a calibrated, cost-constrained policy forecast, and the temporal split, while leakage-free, evaluates on a single contiguous future window.

6. CONCLUSION AND FUTURE WORK

We presented a reproducible, explainable-AI framework that predicts national CO_2 emissions on an open multi-country panel with leakage-free temporal validation, explains the best model with three convergent XAI methods, and translates the explanations into a quantified mitigation counterfactual. Gradient boosting achieved a log-scale R^2 of 0.978, and SHAP, permutation importance and partial-dependence analysis agreed that GDP and energy intensity are the dominant, actionable drivers; a digital-twin experiment showed emission reductions of up to 17% for India under a combined decarbonization lever. Future work will incorporate energy-mix (renewable and fossil share) variables as they become available, fit region- or country-specific models to capture heterogeneity, add calibrated uncertainty via Bayesian or ensemble methods, and replace the manual perturbation with a constrained optimizer that minimizes emissions subject to an economic-growth constraint—moving the digital twin from descriptive to prescriptive.

REFERENCES

[1] Scott M Lundberg, Su-In Lee, "A Unified Approach to Interpreting Model Predictions", Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2017.

[2] Tianqi Chen, Carlos Guestrin, "XGBoost: A Scalable Tree Boosting System", Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 785–794, 2016. https://doi.org/10.1145/2939672.29397

[3] Breiman, L., "Random Forests", Machine Learning, Vol. 45, No.1, pp. 5–32, 2001. https://doi.org/10.1023/A:1010933404324

[4] Jerome H. Friedman, "Greedy function approximation: A gradient boosting machine", Annals of Statistics, Vol. 29, Issue 5, pp.1189–1232, 2001. https://doi.org/10.1214/aos/1013203451

[5] Hannah Ritchie, Pablo Rosado, and Max Roser, "CO_2 and Greenhouse Gas Emissions", Our World in Data, 2024. https://ourworldindata.org/co2-and-greenhouse-gas-emissions

[6] Christoph Molnar, "Interpretable Machine Learning: A Guide for Making Black Box Models Explainable", 2nd Edition, 2025. https://christophm.github.io/interpretable-ml-book/

[7] Arju Manara Begum & Mahadee Al Mobin, "A machine learning approach to carbon emissions prediction of the top eleven emitters by 2030 and their prospects for meeting Paris agreement targets", Scientific Reports, Vol. 15, 2025. https://doi.org/10.1038/s41598-025-04236-5

[8] Siting Hong, Ting Fu, and Ming Dai, "Machine Learning-Based Carbon Emission Predictions and Customized Reduction Strategies for 30 Chinese Provinces", Sustainability, Vol. 17, Issue 5, 2025. https://doi.org/10.3390/su17051786

[9] Pingyang Zhang, Yan Ma, Xujiang Wang, Meng Yang, and Wenlong Wang, "An Interpretable Machine Learning-Based Framework for CO_2 Emission Prediction and Optimization: A Case Study of a University Campus", Sustainability, Vol. 17, Issue 23, 2025. https://doi.org/10.3390/su172310432

[10] Gazi Mohammad Imdadul Alam, Sharia Arfin Tanim, Sumit Kanti Sarker, Yutaka Watanobe, Rashedul Islam, M. F. Mridha & Kamruddin Nur, "Deep learning model based prediction of vehicle CO_2 emissions with eXplainable AI integration for sustainable environment", Scientific Reports, Vol. 15, 2025. https://doi.org/10.1038/s41598-025-87233-y

[11] Dong Yuan, Long Tang, Xueyuan Yang, Fanqin Xu and Kailong Liu, "Explainable Machine Learning Prediction of Vehicle CO2 Emissions for Sustainable Energy and Transport", Energies, Vol. 18, Issue 20, 2025. https://doi.org/10.3390/en18205408

[12] Gökalp Çınarer, Murat Kadir Yeşilyurt, Ümit Ağbulut, Zeki Yılbaşı and Kazım Kılıç, "Application of various machine learning algorithms in view of predicting the CO_2 emissions in the transportation sector", Science and Technology for Energy Transition, Vol. 79, Issue 15, 2024.

[13] https://www.stet-review.org/articles/stet/full_html/2024/01/stet20240008/

[14] Yıldırım Özüpak, Emrah Aslan, "An enhanced interpretable ML approach for forecasting long-term CO_2 and N_2O emissions via optimized random forest modelling", Fuel, Vol. 427, Part A, 2026. https://doi.org/10.1016/j.fuel.2026.139723

[15] Thi Thanh Tam Han, Ching-Yang Lin, "Exploring long-run CO_2 emission patterns and the environmental Kuznets curve with machine learning methods", Innovation and Green Development, Vol. 4, Issue 1, 2024. https://doi.org/10.1016/j.igd.2024.100195

[16] Volkan Şahin, "Explainable machine learning-based prediction of CO_2 emissions from passenger vessels", Ocean Engineering, Vol. 341, Part 3, 2025. https://doi.org/10.1016/j.oceaneng.2025.122752

[17] Wenyang Wang, Yuping Luo, Zihan Jiang, Jibin Zhou, Peng Jia, "A deep learning framework for global transportation energy carbon-emission forecasting: integrating a generative pre-trained transformer with multi-scale feature analysis", Energy, Vol. 338, Issue 30, 2025. https://doi.org/10.1016/j.energy.2025.138586

[18] Mohd Azfar Shaida, Laiba Saleem, Syed Ali Waqas Ahmad, Saad Shamim Ansari & Syed Muhammad Ibrahim, "Data-driven explainable artificial intelligence models for predicting CO_2 uptake in metal–organic frameworks", Environmental Science and Pollution Research, Vol. 33, pp. 6591–6609, 2026. https://doi.org/10.1007/s11356-026-37725-9

[19] Salim Oyinlola, Joshua Ajayi, Gozie Ibekwe, Chika Yinka-Banjo, "Time Series Based CO_2 Emission Forecasting and Energy Mix Analysis for Net Zero Transitions: A Multi Country Study", arXiv preprint, 2026. https://doi.org/10.48550/arXiv.2601.01105

[20] Intergovernmental Panel on Climate Change (IPCC): Climate Change 2023: Synthesis Report. IPCC, Geneva, Switzerland, 2023. https://report.ipcc.ch/ar6syr/

[21] Ministry of Environment, Forest and Climate Change, Government of India: India's Updated First Nationally Determined Contribution under the Paris Agreement. New Delhi, 2026. https://unfccc.int/sites/default/files/2026-04/INDIA%20NDC%202031-35.pdf