

Machine Learning Analysis for Fuel Consumption in Heavy Vehicles

Nurudeen A. Raji¹, Medekannu M. Ogunbiyi², Abdullateef O. Bakare³, Choice Eseoghene⁴, Ojo Seun⁵
^{1,2,3,4,5} *Mechanical Engineering Department, lagos State University, Nigeria*

Abstract—Efficiently managing heavy vehicle fuel consumption is still a top priority for the transportation sector. With the use of cutting-edge machine learning techniques, this research aims to develop a robust model that can predict fuel consumption with improved accuracy given a variety of features from different heavy vehicles. The project uses a data-driven approach and draws on a sizable real-world dataset from a variety of heavy vehicle. To assure data integrity and coherence, this dataset has undergone painstaking processing and improvement. We explore the dataset's intricate patterns and using exploratory data mining techniques. To identify the most relevant predictors impacting fuel consumption dynamics, feature selection and engineering are applied as integral parts of the process. A wide range of machine learning techniques such as regression analysis, ensemble learning (voting and stacking regressor), hyperparameter tuning, neural network dropout and batch normalization underwent careful training and evaluation. With thorough evaluation, we identify stacking ensemble model as the front-runner with a stunning accuracy of 0.999983 and MSE (mean square errors) of 0.001007 and (support vector machine) SVM, XGBoost, gradient boosting, and random forest by a rigorous comparative analysis, revealing them to be the best performing among the models after hyperparameter tuning 0.999968% , 0.999525%, 0.99821%, 0.996388% 2, and MSE of 0.00186, 0.028248, 0.105822, 0.214788 respectively. Our analysis underscores the importance of testing multiple techniques before drawing definitive conclusions. Particularly, the Support Vector Machine (SVM) initially exhibiting poor performance with an accuracy of 23.36% and MSE of 45.57, highlighting the necessity for thorough evaluation.

Index Terms—fuel consumption prediction, machine learning, heavy vehicles, regression analysis, ensemble Learning.

I. INTRODUCTION

Predicting the fuel consumption of heavy vehicles is a crucial part of managing and planning the transportation system. Due to the substantial impact that heavy trucks have on the transportation industry, precisely forecasting fuel consumption can result in significant financial savings, increased operational effectiveness, and less environmental impact. Machine learning analysis has become a useful method for predicting fuel consumption in heavy vehicles. Predictive models can be created to estimate fuel consumption based on input variables, and other pertinent aspects by utilizing the power of machine learning algorithms. Model-based estimation is one strategy for forecasting fuel consumption, as written by [2].

To ensure accurate predictions, this strategy is sensitive to uncertainties and calls for strict machine learning methods. Eco-driving is one industry where machine learning has been successfully used. In order to limit fuel consumption and cut emissions, eco-driving involves improving driving methods and routes. For various portions of a route, machine learning approaches have been applied to predict fuel usage and, as a result, the emission of CO₂ [6].

The use of machine learning analysis in heavy vehicle fuel consumption prediction has shown significant promise for enhancing operational effectiveness, cutting costs, and lowering environmental impact. machine learning analysis has proven to be useful in predicting the fuel consumption of heavy vehicles. By considering input variables including vehicle features, driving circumstances, and other pertinent elements, these models may efficiently estimate fuel usage. Based on historical data and real-time inputs, machine learning algorithms can be used to analyze and predict fuel usage patterns [7].

Fuel consumption prediction for heavy vehicles is a key problem in the field of transportation engineering, with implications for fleet management, environmental effect, and the broader field of mechanical engineering [5].

In addition to the financial ramifications, fuel usage analysis has a substantial environmental impact. Heavy vehicles are responsible for a considerable amount of greenhouse gas emissions and air pollution [4]. By precisely anticipating fuel use, fleet managers and policymakers may apply policies to decrease the environmental footprint of heavy vehicles. This includes optimizing routes, adopting fuel-efficient devices, and promoting eco-driving practices. Ultimately, this research project intends to contribute to sustainable transportation practices and decrease the environmental impact of heavy vehicle operations.

In recent years, many initiatives have been made to solve the challenge of fuel usage in large trucks. However, precisely estimating fuel usage remains a hard endeavor due to the interplay of multiple factors. The necessity of precisely anticipating fuel usage in big vehicles becomes clear when considering the cost ramifications. Fuel expenditures account for a major component of the overall operational expenses for fleet operators [3]. Therefore, precise prediction of fuel consumption is vital for effective fleet management and optimization. Fuel consumption analysis has a substantial environmental impact. Heavy vehicles are substantial contributors to greenhouse gas emissions. These factors include vehicle attributes, driving circumstances, and driver behavior, all of which influence the fuel efficiency of heavy trucks which contribute to climate change as heavy-duty vehicles account for over 30 percent of on-road CO₂ emission in the European Union alone [2]. Additionally, the combustion of fossil fuels in these cars contributes to air pollution, significantly harming air quality and public health. By increasing fuel consumption prediction, fleet operators and policymakers may implement policies to lower the environmental imprint of heavy trucks. While various research has explored fuel consumption in large vehicles, there are still gaps in understanding the precise influence of each aspect and the development

of accurate forecasting models [1]. This research intends to solve these gaps by conducting a detailed analysis of factors impacting fuel consumption in heavy vehicles and constructing a viable forecasting model. By investigating a varied dataset of fuel consumption and associated variables, this research project will provide useful insights into the intricate relationship between fuel consumption and many factors.

II. METHODOLOGY

In this study, we adopted Python V3.10.1 and development environment was the Google Colab notebook Integrated Development Environment to perform a thorough examination of midsize trucks, minivans, and other vehicles. Data collection, preprocessing, feature engineering, model selection, evaluation, and interpretation were all divided into discrete stages of the process. A variety of midsize trucks, minivans, and other automobiles provided the dataset used in this study. The information covered a wide range of characteristics, including fuel economy, engine size, transmission type, and more. In order to prepare the data for modeling, this required addressing missing values, encoding category variables, and scaling numerical characteristics.

Feature engineering was essential in improving the models' capacity for prediction. To better capture trends in the data, new features were created and old features were modified. A number of machine learning models were investigated in order to determine which would be best for forecasting midsize truck and minivan fuel consumption. Neural networks, XGBoost, Random Forest, Support Vector Machine (SVM), and Gradient Boosting were among the options. To maximize the performance of the selected models, hyperparameter adjustment was carried out. To evaluate the models' accuracy, metrics such Mean Squared Error (MSE), Mean Absolute Error (MAE), and R-squared were used.

For our analysis on fuel consumption prediction, we have collected a dataset with pertinent data on numerous mini sized trucks, as well as the fuel consumption and CO₂ emissions that go along with them. Included in the collection are attributes like "MAKE," "MODEL," "VEHICLECLASS," "ENGINE SIZE," "CYLINDERS," "FUELTYPE," "FUELCONSUMPTION_CITY",

"FUELCONSUMPTION_HWY,"
 "FUELCONSUMPTION_COMB," and
 "CO2EMISSIONS."

The dataset comprises actual field observations of various mini sized trucks and vans and their corresponding fuel consumption rates along with other features recorded from the vehicles. It was gathered from a dependable source. Before moving further with model building, it is crucial to validate the integrity and quality of the dataset, as with any other dataset.

Handling missing data is an important component of data preparation. Fortunately, the dataset we used for our fuel consumption prediction project does not contain any missing numbers. Missing data can occur for a number of reasons, including incorrect data collection, sensor malfunctions, or user omissions. In contrast, every entry in our dataset is full, and every observation has values for every characteristic.

However, it is crucial to use caution while gathering data and preparing it. For the machine learning models to produce correct predictions, it is essential that the data is collected precisely and comprehensively. Additionally, to maintain the dataset's quality and integrity, we must continue to check for missing values and manage them appropriately if you come across new data in the future or intend to expand it. We can confidently move on to the next processes of feature selection, engineering, model creation, evaluation, and simulation approaches because our dataset does not contain any missing data. The data's completeness enables us to fully study and utilize the information at our disposal, laying the groundwork for a reliable prediction of heavy trucks' fuel usage.

Categorical variables, such as "MAKE," "MODEL," "VEHICLECLASS," and "FUELTYPE," are not numeric and must be transformed into numerical representations before the models can process them. These category features were converted into binary vectors using one-hot encoding to make sure the models can utilise the data they give.

We divided the preprocessed data into training and testing sets before developing the model. The testing set will be used to assess the machine learning models' performance once they have been trained on the training set. This makes it possible for us to evaluate how well the model generalizes to new data. In order to evaluate how well a machine learning model generalizes to new data, data splitting is a crucial phase in the model building process. The training set and the

testing set are the two subsets that make up the dataset. The model is trained on the training set, and its performance is assessed on the testing set.

In our fuel consumption prediction project, we divided the data into a training set and a testing set, with a split of 70% and 30%, respectively. Accordingly, 70% of the data was utilized to train the machine learning models, while the remaining 30% was set aside for testing the models' effectiveness on fresh, untried data. The machine learning models have a sizable amount of data to learn from thanks to the allocation of 70% of the data for training. The model memorizes the training data but is unable to generalize to new data due to overfitting, which can be avoided with enough training data. By setting aside 30% of the data for testing, we may assess the model on a large amount of previously unobserved data. Particularly if the dataset is constrained in size, a smaller testing set aids in ensuring that the model's performance on fresh data is accurately reflected.

The 70-30 data split strikes a balance between giving sufficient data for training and guaranteeing a sufficient sample size for testing. This ratio is frequently employed and has performed well in a number of machine learning applications. The feature matrix X and the target vector y are denoted by the letters X and y , respectively. The target variable (in this case, the column labeled "FUELCONSUMPTION_COMB" reflecting fuel consumption combined for fuel consumption in the city and highway) is contained in column y , whereas all independent variables (features) from the dataset are contained in column X . `test_size=0.3`: This value indicates how much of the dataset will be used for testing. We have set the `test_size` parameter to 0.3, which means that 30% of the data will be used for testing and the remaining 70% for training. The data splitting process's unpredictability is managed using the `random_state=42` parameter. We guarantee the reproducibility of the data splitting by setting `random_state` to a fixed value, such as 42. Four variables are returned by the `train_test_split` function: `X_train`, `X_test`, `y_train`, and `y_test`. The training and testing sets of features and target values are correspondingly contained in these variables. The machine learning models was trained using the `X_train` and `y_train`, and their performance was assessed using the `X_test` and `y_test`.

To evaluate the model's capacity to generalize to new data, the data is divided into training and testing sets. The model is trained by using the training set to identify trends, and it is then tested against the testing set to see how well it can forecast fuel use for brand-new, unforeseen data. We find a balance between supplying enough data to train the models and having a sizeable chunk of data for trustworthy evaluation by utilizing a 70-30 data split. By doing this, overfitting is avoided and the model's performance indicators, including mean squared error and R-squared, reflect its capacity for making predictions using fresh, real-world data.

2.1 Evaluation Techniques

Models Based On Regression: Fuel consumption forecasting has made substantial use of regression-based models. For example, the link between input parameters (such as engine size, cylinder count, and vehicle weight) and fuel consumption were modelled using regression models.

Hyperparameter Tuning: A crucial stage in the creation of machine learning models is hyperparameter tuning, which seeks to maximize a model's performance by determining the ideal hyperparameter configuration. An intensive hyperparameter tuning method was utilized in our fuel consumption analysis of several vehicles, including minivans and midsize trucks, in order to improve the predicted accuracy of the selected models.

Ensemble Learning: Ensemble learning is a powerful technique in machine learning that combines forecasts from different models to create a more robust and precise model. During our investigation of fuel usage in midsize trucks, minivans, and other vehicles, we examined two commonly used ensemble learning techniques: Voting Regressor and Stacking Regressor.

Voting Regressor: The Voting Regressor is an ensemble technique that enhances the overall predictive performance by aggregating the predictions of numerous base estimators (models). When dealing with regression problems, the predictions are combined, typically by calculating the average, in order to provide a final forecast that is more reliable and precise. In our investigation, we examined several fundamental models, such as Random Forest, XGBoost, SVM, Neural Network, and Gradient Boosting. The hyperparameters of each base model

were individually optimized utilizing approaches such as Grid Search Cross-Validation.

Stacking Regressor: Stacking Regressor, like Voting Regressor, is an ensemble learning strategy that integrates predictions from numerous base models. However, unlike the uniform averaging in Voting Regressor, Stacking Regressor introduces a meta-model that learns how to best combine the forecasts of the base models. The base models (Random Forest, XGBoost, SVM, Neural Network, Gradient Boosting) were separately tweaked, and a Stacking Regressor was developed using a meta-model, in our case, a Linear Regression model. Ensemble learning, through approaches like Voting Regressor and Stacking Regressor, was successful in boosting the prediction capabilities of our models for fuel usage analysis. By combining the strengths of multiple models, these ensemble methods addressed the limits of individual models, leading to enhanced accuracy and robustness in estimating fuel consumption for midsize trucks, minivans, and other vehicles.

Dropout And Batch Normalization : Neural Networks, with their ability to capture complex patterns in data, play a significant role in predictive modeling. In our fuel consumption research, we enhanced the classic neural network architecture by including two powerful technique.

Dropout: Dropout is a regularization approach that boosts the generalization ability of neural networks by preventing overfitting. It operates by randomly "dropping out" a fraction of neurons during training, causing the network to learn redundant representations. This randomization helps prevent co-adaptation of hidden units and resulting in a more resilient model.

Batch Normalization: Batch Normalization is a technique that normalizes the input of each layer across a mini-batch. It aids in reducing difficulties like internal covariate shift, making the training process more stable and quickening convergence. Batch Normalization is done to the inputs of hidden layers. Batch Normalization is smoothly integrated into the neural network architecture utilizing the BatchNormalization layer in Keras. After establishing the neural network models using Dropout and Batch Normalization, training is performed on the training set. The models are then assessed on the test set using metrics such as Mean Squared Error (MSE), Mean Absolute Error (MAE), and R-squared. The

integration of Dropout and Batch Normalization in our neural network models boosts their resilience, generalization, and convergence speed.

III. RESULTS AND DISCUSSIONS

Regression model: The regression model and their results are presented in table 1 and figure 1. The results show the result for each of those models providing comprehensive strength of each model and their respective performance on the fuel consumption dataset. The contributing regression models adopted in this research work include XGBoost, Support Vector Machine (SVM), Random Forest, Neural Network, and Gradient Boosting. The performance of each contributing regression model shows their respective capacity to predict factors and metrics contributing to fuel consumption. The evaluation metrics, such as Mean Squared Error (MSE), Mean Absolute Error (MAE), R-Squared, and Accuracy, offer a comprehensive understanding of the performance of each model. The Random Forest model demonstrates remarkable predictive accuracy, as evidenced by its low Mean Squared Error (MSE) and Mean Absolute Error (MAE), which indicate exact estimates of fuel use. The R-Squared score of 0.996388 indicates that the model is able to explain a significant portion of the variability in the data.

Gradient Boosting demonstrates higher performance, as seen by its reduced Mean Squared Error (MSE) and Mean Absolute Error (MAE) values in comparison to Random Forest. The amazing R-Squared of 0.998221 shows the model's precision in identifying complicated relationships within the data. XGBoost exceeds other models, featuring the lowest MSE and MAE. The remarkable R-Squared value of 0.999525 indicates an extremely accurate fit, highlighting its strong predictive ability.

Support Vector Machines (SVM), on the other hand, encounter difficulties when it comes to accurately forecasting fuel usage for large vehicles, as seen by the notably elevated Mean Squared Error (MSE) of 45.371631. The relatively low R-Squared score of 0.237070 indicates that there are limits in accurately capturing the underlying patterns in the data. The Neural Network exhibits robust predictive skills, although it does not match the level of performance attained by ensemble models such as XGBoost with

accuracy of 0.999525 and Gradient Boosting 0.999525.

Table 1 Model Performance metric for each Model (Result for performance metrics was gotten from the machine learning analysis).

Models	MSE	MAE	R-Squared	Accuracy
Random Forest	0.238409	0.065452	0.995991	0.9963
Gradient Boosting	0.097678	0.037725	0.998358	0.998221
XG Boost	0.018084	0.029385	0.999696	0.999525
SVM	45.57468	5.204254	0.233656	0.23707
Neural Network	3.232715	1.042876	0.945641	0.984764




Fig. 1 Model accuracy for regression models

The comparatively elevated Mean Squared Error (MSE) and Mean Absolute Error (MAE) suggest a little less precise alignment. XGBoost and Gradient Boosting, surpass individual models such as SVM and Neural Networks in accurately predicting fuel consumption for heavy vehicles. The strong accuracy scores show the efficacy of ensemble methods in catching detailed patterns within the data, making them valuable tools for real-world fuel consumption prediction.

The mean absolute error bar chart in figure 2 was used to evaluate the accuracy of our models by measuring the average absolute differences between the actual and predicted values. The low MAE indicates for XGBoost, Neural Network,




Fig. 2 Mean absolute error for model performance

Random Forest, Gradient Boost, suggests that on average, the model's predictions are close to the actual values. Lower MAE values suggest better accuracy and precision in predicting the target variable. SVM and Neural Network on the other hand with higher MAE suggests that the model's predictions are, on average, farther away from the actual values.

A low R^2 value, close to 0 as seen from the visualisation in fig 3 indicates that SVM is not capturing much of the variability in the target variable. The predictions do not align well with the actual values, suggesting that SVM model is not explaining the patterns in the data effectively.

The high R^2 squared values for other models which are close to 1, suggests that the models is explaining a significant portion of the variability in the target variable. The predictions closely align with the actual values, indicating a good fit to the data.




Fig. 3 R-Squared comparison for model performance

The bar char for low mean squared error in fig 4 indicates that the model's predictions are, on average, close to the actual values. In other words, the models are making accurate predictions, and the squared

differences between predicted and actual values are small. SVM and Neural Network performance shows that the predictions are far off the actual values.




Fig. 4 Mean Square Error for model performance

In the comparison of expected vs. predicted values for our regression models in fig 5, the visual clearly illustrates the performance of each model. Notably, the Support Vector Machine (SVM) exhibits a deviation from the expected values, with data points indicating a noticeable discrepancy between predicted and actual outcomes. This divergence is distinct from the more closely aligned points observed with other models such as Random Forest, Gradient Boosting, XGBoost, and Neural Network. The visualization underscores the variance in predictive accuracy among the models, emphasizing the unique characteristics of each algorithm in capturing the underlying patterns in the data.




Fig. 5 Predicted values vs Expected values for all models

Hyperparameter Tuning: In this part, we describe how several machine learning algorithms are used to build

and train models for forecasting fuel use in heavy trucks. The pre-processed dataset has been divided into training and testing sets. We use a variety of regression-based models to investigate their propensity for prediction. These models include a Multi-Layer Perceptron (Neural Network), Random Forest, Gradient Boosting, XGBoost, and others. The Random Forest model, following hyperparameter adjustment, continues to display remarkable performance. With appropriate settings, the model obtains a low MSE and MAE, suggesting accuracy and precision in predicting fuel usage. The hyperparameter-tuned XGBoost model maintains its powerful prediction ability. Although it has a slightly greater MSE and MAE compared to Random Forest, it still performs incredibly well. The high R-squared value and accuracy indicate the model's usefulness. The hyperparameter-tuned SVM model demonstrates remarkable improvement, achieving extremely low MSE and MAE.

The hyperparameter-tuned Neural Network model, while not surpassing other models, displays increased performance. The reduction in MSE and MAE, together with a decent R-squared value and accuracy, suggests better predictive potential. Hyperparameter adjustment greatly enhances the Gradient Boosting model, resulting in lower MSE and MAE. The strong R-squared value and accuracy show its improved precision in estimating fuel consumption for heavy vehicles.

Table 2. Performance metrics for hyperparameter tuning for each model (Result for performance metrics was gotten from the machine learning analysis).

Models	MSE	MAE	R-Square d	Accura cy
Random Forest	0.25695	0.06	0.9953	0.9956792
Gradient Boosting	0.04099	0.03138	0.99931	0.9993107
XG Boost	0.42686	0.08263	0.99201	0.9928222
SVM	0.00186	0.40612	0.99997	0.9999687
Neural Network	0.44864	0.03138	0.99931	0.9993107

Voting And Stacking Regressor: Fuel consumption analysis for heavy vehicles demands powerful predictive models that can effectively capture the intricacies of the underlying data. In this area, ensemble learning approaches such as Voting and Stacking Regressors have been applied to enhance the accuracy and robustness of predictions.

Voting Ensemble Model: The Voting Regressor collects predictions from numerous base models, generating a powerful ensemble model. In our analysis, the Voting Ensemble Model showed considerable improvements.

- MSE: 0.12816774438360887
- MAE: 0.10905163915238206
- R-squared: 0.9978448426731634
- Accuracy: 0.9978448426731634

The new results demonstrate solid performance, offering improved accuracy and dependability compared to individual models. The ensemble effect contributes to minimizing model biases, resulting in more consistent and precise predictions for fuel consumption.

Stacking Ensemble Model: The Stacking Regressor includes training a meta-model to aggregate predictions from multiple base models. The new results for the Stacking Ensemble Model reveal considerable improvements with best performance output and has proven to be most .

- MSE: 0.0010072817015525598
- MAE: 0.022202728661158643
- R-squared: 0.9999830624269022
- Accuracy: 0.9999830624269022

The Stacking Ensemble Model displays tremendous enhancements in predicted accuracy and reliability. The meta-model optimally weights predictions from base models, resulting in a highly accurate and robust prediction of fuel consumption.

The Voting and Stacking Regressors routinely outperform individual models, displaying their ability to use the strengths of various algorithms for more accurate predictions. Ensemble learning strengthens the robustness of our analysis by lowering the influence of outliers or mistakes contained in individual models. The implementation of Voting and Stacking Regressors in our fuel consumption analysis has shown to be highly effective. The latest results

reveal significant improvements in predicted accuracy and reliability. The ensemble learning techniques have successfully addressed the constraints of individual models, offering a more comprehensive and precise understanding of fuel efficiency in heavy vehicles. These strategies stand out as valuable instruments in our analytical toolset, adding significantly to the effectiveness of our fuel consumption prediction models.

Neural Network With Dropout & Batch Normalisation: Neural Networks equipped with advanced regularization techniques, notably Dropout and Batch Normalization, have been applied for predicting fuel consumption in heavy vehicles. These strategies are aimed to enhance the model's generalization and stability by addressing frequent concerns including overfitting and internal covariate shifts. The neural network architecture employed for this fuel consumption analysis has numerous layers, combining Dropout and Batch Normalization.

Input Layer: Represents the features of the fuel consumption dataset.

Hidden Layers: Utilize complex activation algorithms to capture subtle patterns.

Dropout Layers: Introduce regularization by randomly setting a fraction of input units to zero during training, reducing overfitting.

Batch Normalization Layers: Normalize the inputs of each layer, stabilizing and accelerating the training process. The Neural Network model with Dropout and Batch Normalization produces impressive results, exhibiting its ability in capturing intricate correlations within the fuel consumption dataset.

- Mean Squared Error (MSE): 15.27
- Mean Absolute Error (MAE): 2.73
- R-squared (R2): 0.74

Dropout serves as a robust regularization approach, limiting overfitting by preventing the network from too relying on specific neurons. This fosters a more generalized learning. While Bath Normalization enhances the stability and speed of training by normalizing inputs, allowing for faster convergence and improved performance.

On performance and generalization, the Neural Network with Dropout and Batch Normalization

exhibits a propensity for generalizing well across varied contexts, suggestive of a robust predictive model. The integration of Dropout mitigates overfitting concerns, ensuring the model does not memorize noise in the training data. The Neural Network model, reinforced with Dropout and Batch Normalization, emerges as a strong solution for fuel consumption prediction. The obtained results underline the efficiency of these sophisticated regularization strategies in boosting the model's generalization, stability, and overall predictive accuracy.

Taking the results for our research into context, the interpretation of high MSE and high R-squared depends on the unique goals and constraints of the study. From the results, it shows that a considerable percentage of the variance in fuel consumption is explained by the features in the model. This indicate a good overall fit. The high MSE, however, means that there are substantial inaccuracies in estimating individual fuel consumption values. This would signal that, while the model captures the overall trend, it struggles with precise forecasts for specific occurrences.

If the primary goal is to have accurate estimates for individual fuel consumption values, minimizing MSE can be a priority. This could involve improving features, incorporating additional data, or investigating various model architectures. However, for this analysis, explaining the overall variability in fuel consumption is crucial. Since this helps to increase our understanding of the factors impacting consumption trends, an emphasis on enhancing R-squared is appropriate. R-Square score of 0.74 is still higher performance considering the primary purpose of this analysis.

The Training Loss Curve (figure 6, shown by the Blue Line) indicates the value of the loss function (namely, the mean squared error) on the training set at each epoch. During the process of training, the model aims to reduce its loss by iteratively modifying its weights across multiple epochs. A diminishing training loss signifies that the model is acquiring knowledge and enhancing its performance on the training data.




Fig. 6 Training and validation loss evolution during neural network training epochs

The Validation Loss Curve (shown by the Orange Line) displays the loss on the unseen dataset (known as the validation set) that the model has not been exposed to during the training process. It functions as a measure of the model's ability to generalize to unfamiliar data. Since the validation loss reduces in conjunction with the training loss, it indicates that the model is not overfitting and is expected to exhibit good performance on unseen data. A substantial disparity between the training and validation loss could suggest the occurrence of overfitting.

An effective model should have a declining pattern in both the loss during training and the loss during validation. If the training loss diminishes, while the validation loss begins to climb or remains elevated, it may indicate the occurrence of overfitting. Overfitting arises when the model excessively captures the intricacies of the training data, encompassing its noise and outliers, yet struggles to apply this knowledge to unseen data. The goal is to have both training and validation losses minimized. If the model is successful, it should perform well on both the training set and new, unseen data. Monitoring the loss curves assists in understanding the model's learning process and making adjustments to improve its generalization capacity.

IV. CONCLUSION AND RECOMMENDATIONS

Our extensive analysis of heavy vehicle fuel consumption using advanced machine learning approaches has given significant insights crucial for enhancing fuel efficiency in the transportation sector. Through painstaking data processing, feature engineering, and model evaluation, we found

numerous critical findings that lead to the advancement of fuel consumption prediction models. Our analysis demonstrated the superiority of ensemble learning approaches, notably Support Vector Machine, XGBoost, Gradient Boosting, and Random Forest, exhibiting remarkable accuracies. These models not only revealed remarkable prediction ability but also exhibited minimum mean squared errors, demonstrating their usefulness in capturing the subtleties of fuel consumption dynamics.

However, our work also underlined the need of comprehensive model evaluation. The Support Vector Machine (SVM) model initially displayed inadequate performance with an accuracy of **23.36%** and a mean squared error of 45.57 further hyperparameter tuning led to a remarkable accuracy of **0.999968%**. This underlines the significance of careful consideration and testing of numerous methodologies before obtaining conclusive findings. Building upon our findings, we propose the following recommendations for additional research and practical implementation:

Model Ensemble Approach: Future research should study the synergistic impact of combining the capabilities of several models through ensemble methods, potentially leading to even more robust fuel consumption prediction models.

Feature Importance refining: Continuous refining of feature engineering methodologies and research of additional relevant characteristics could boost the models' ability to capture nuances in heavy vehicle fuel consumption patterns.

Real-time Monitoring: Implementing real-time monitoring systems based on our verified models could empower fleet managers and regulators to make prompt decisions, optimizing fuel use across varied scenarios.

REFERENCES

- [1] Barbado González A, Corcho O. Understanding Factors Affecting Fuel Consumption of Vehicles Through Explainable AI: A Use Case With Explainable Boosting Machines. 2021. [online] Available from <https://www.researchgate.net/publication/35323>

- 5325_Understanding_Factors_Affecting_Fuel_Consumption_of_Vehicles_Through_Explainable_AI_A_Use_Case_With_Explainable_Boosting_Machines
- [2] Budzyński A, Ślaskowski A. Forecasting fuel consumption in means of transport with the use of machine learning. *Transport Technologies*. 2022. [online] Available from: https://www.researchgate.net/publication/366176279_Forecasting_fuel_consumption_in_means_of_transport_with_the_use_of_machine_learning [Accessed: 7th July 2023]
 - [3] Budihardjo MA, Fadhilah I, Humaira NG, Hadiwidodo M, Wardhana IW, Ramadan BS. Forecasting Greenhouse Gas Emissions from Heavy Vehicles: A Case study of Semarang City. *Jurnal Presipitasi: Media Komunikasi dan Pengembangan Teknik Lingkungan*. 2021;18(2):254-260. <https://doi.org/10.14710/interaksi.v.v.i.p>.
 - [4] Li D, Li C, Miwa T, Morikawa T. An exploration of factors affecting drivers' daily fuel consumption efficiencies considering multi-level random effects. *Sustainability*. 2019; <https://doi.org/10.3390/su11020393>
 - [5] Matej K, Sciarretta A, Arben C, Mounier H, Thibault L, Niculescu SI. About prediction of vehicle energy consumption for eco-routing. 2016:1096-1101. 10.1109/ITSC.2016.7795693. [online] Available from: https://www.researchgate.net/publication/312247980_About_prediction_of_vehicle_energy_consumption_for_eco-routing
 - [6] Sun L, Zhang Y, Hu X. Economical-traveling-distance-based fleet composition with fuel costs: An application in petrol distribution. *Transportation Research Part E: Logistics and Transportation Review*. 2021;147. [online] Available from: <https://ideas.repec.org/a/eee/transe/v147y2021ics1366554521000016>.
 - [7] Subramaniam N, Yusof N. Modelling of CO2 Emission Prediction for Dynamic Vehicle Travel Behavior Using Ensemble Machine Learning Technique. 2021:383-387. <http://dx.doi.org/10.1109/SCORED53546.2021.9652757>