information

MDPI

Article

Explainable AI Using On-Board Diagnostics Data for Urban Buses Maintenance Management: A Study Case

Bernardo Tormos [1], Benjamín Pla [1], Ramón Sánchez-Márquez [1,*] and Jose Luis Carballo [2]

[1] CMT-Clean Mobility & Thermofluids, Universitat Politècnica de València, 46022 Valencia, Spain; betormos@mot.upv.es (B.T.); benplamo@mot.upv.es (B.P.)
[2] Empresa Municipal de Transportes de Valencia S.A.U., 46001 Valencia, Spain; jcarballo@emtvalencia.es
* Correspondence: resanmar@mot.upv.es

Abstract: Industry 4.0, leveraging tools like AI and the massive generation of data, is driving a paradigm shift in maintenance management. Specifically, in the realm of Artificial Intelligence (AI), traditionally "black box" models are now being unveiled through explainable AI techniques, which provide insights into model decision-making processes. This study addresses the underutilization of these techniques alongside On-Board Diagnostics data by maintenance management teams in urban bus fleets for addressing key issues affecting vehicle reliability and maintenance needs. In the context of urban bus fleets, diesel particulate filter regeneration processes frequently operate under suboptimal conditions, accelerating engine oil degradation and increasing maintenance costs. Due to limited documentation on the control system of the filter, the maintenance team faces obstacles in proposing solutions based on a comprehensive understanding of the system's behavior and control logic. The objective of this study is to analyze and predict the various states during the diesel particulate filter regeneration process using Machine Learning and explainable artificial intelligence techniques. The insights obtained aim to provide the maintenance team with a deeper understanding of the filter's control logic, enabling them to develop proposals grounded in a comprehensive understanding of the system. This study employs a combination of traditional Machine Learning models, including XGBoost, LightGBM, Random Forest, and Support Vector Machine. The target variable, representing three possible regeneration states, was transformed using a one-vs-rest approach, resulting in three binary classification tasks where each target state was individually classified against all other states. Additionally, explainable AI techniques such as Shapley Additive Explanations, Partial Dependence Plots, and Individual Conditional Expectation were applied to interpret and visualize the conditions influencing each regeneration state. The results successfully associate two states with specific operating conditions and establish operational thresholds for key variables, offering practical guidelines for optimizing the regeneration process.

Keywords: explainable artificial intelligence; maintenance management; diesel particulate filter; urban bus fleets; machine learning

check for updates

Academic Editor: Tieling Zhang

Received: 17 December 2024
Revised: 13 January 2025
Accepted: 17 January 2025
Published: 21 January 2025

Citation: Tormos, B.; Pla, B.; Sánchez-Márquez, R.; Carballo, J.L. Explainable AI Using On-Board Diagnostics Data for Urban Buses Maintenance Management: A Study Case. *Information* **2025**, *16*, 74. https://doi.org/10.3390/info16020074

1. Introduction

Maintenance management has experienced a paradigm shift with the advent of Industry 4.0, bringing the integration of advanced tools such as Artificial Intelligence (AI), the Internet of Things (IoT), Big Data, and Machine Learning. These developments enable the capture, analysis, and processing of real-time data, transforming traditional maintenance into a digitized, predictive, and proactive approach [1,2]. Unlike conventional approaches, which rely on scheduled preventive or corrective maintenance, digitized maintenance leverages the analysis of large data volumes to predict failures and optimize asset lifecycles [3,4]. However, implementing these new technologies comes with challenges. Integrating data from diverse sources can be complex. The maintenance team often lacks the technical expertise required to leverage advanced tools, and the high initial costs associated with adopting these technologies pose barriers to broader and more seamless implementation.

The use of advanced data analysis techniques, such as Machine Learning, as a way to improve vehicle maintenance management has been previously explored in various studies. For instance, one of the studies developed a remote diagnostic and maintenance system leveraging least squares support vector machines (LS-SVM), demonstrating strong performance with 93% accuracy. Despite notable advantages, such as enabling more proactive maintenance by estimating the Remaining Useful Life (RUL), it faces challenges including integration complexity and high initial costs [5]. Another study proposed an ensemble-based classifier approach for detecting known and unknown faults in vehicles using multivariate road test data. By combining binary and one-class classifiers, the method achieved an average F2 score of 77% and robustness of 85% under varying driving conditions and fault types. Notably, this solution is 'out-of-the-box,' requiring no expert parameterization and supporting predictive maintenance and condition monitoring. However, its limitation lies in being restricted to offline analysis, precluding real-time application [6]. Another significant contribution was the development of a Random Forest model to predict air compressor failures in trucks using historical vehicle data and maintenance logs. By employing beam search for feature selection, the model achieved 77% accuracy and estimated economic savings of 1.66 million euros. This approach proved effective for noisy and imbalanced data typical of the automotive industry, but its use is constrained by the low frequency of available data and its offline-only application [7]. The development of models employing algorithms such as decision trees and support vector machines (SVM) has also been studied to prevent failures in critical components, including transmissions, electrical systems, and batteries. The results included specific metrics such as an 80% accuracy rate in predicting critical events for fleets, alongside benefits like optimized maintenance schedules and reduced downtime. Advantages included extending component life and integrating IoT data for real-time analysis. However, limitations were identified in handling noisy and missing data, as well as challenges in implementing interpretable models for production environments [8]. Another approach to enhance operation and maintenance management in the automotive industry was developed using scrap crawlers to collect historical data and ID3 decision tree algorithms to classify faults and prioritize maintenance actions. Results indicated high accuracy in fault classification, emphasizing the scalability and practicality of the approach. Advantages included task automation and reduced response time to failures. However, limitations involved reliance on small datasets and the inability of the ID3 algorithm to efficiently handle continuous variables [9]. In the agricultural domain, a study applied supervised autonomous learning to manage the maintenance of tractor engines. Using decision tree classification algorithms, the study simulated fuel injection system failures and collected vibration data to diagnose engine conditions. The results demonstrated 95% overall classification efficiency and up to 97.5% in individual training sessions, validating its applicability for predictive diagnostics. Advantages included quick fault identification and resource savings, while limitations centered on reliance on a limited dataset and the need for manual intervention to characterize the data [10].

These studies have demonstrated how the application of AI models allows the detection of patterns that can anticipate failures, thereby optimizing asset management and consequently extending their useful life. In the context of urban transport fleets, where the availability and reliability of vehicles are crucial to maintaining the quality of service provided, digitized maintenance offers a unique opportunity to improve operational efficiency, reduce downtime, and extend the lifespan of critical components. The ability to monitor fleet health in real time and predict the need for intervention enables a much more dynamic and personalized approach tailored to the specific usage conditions of each fleet [11,12].

Specifically addressing the use of Machine Learning as a way to improve maintenance management in urban bus fleets, various studies have been conducted to demonstrate its potential. The COSMO approach proposes a self-organized system model for fault detection in urban buses, leveraging sensor data already available in the vehicle. This method, validated with both simulated and real-world data, successfully detected all cooling system failures in tests conducted on an urban bus. While it offers the advantage of adapting to different vehicle configurations without the need for prior modeling, it faces limitations in detection accuracy for critical systems or those with complex variability [13]. The ICOSMO is an improved version of the previous study, which introduces an IoT architecture combined with semi-supervised Machine Learning to enhance sensor selection for predictive maintenance in public bus fleets. Validated through a prototype implemented at the Société de Transport de l'Outaouais (STO), this system processed 1 GB of J1939 data daily, enabling anomaly detection in critical sensors such as engine speed and fluid temperatures. Although the initial implementation was limited to a single bus, the approach demonstrated scalability potential for entire fleets, despite challenges related to costs and infrastructure requirements [14]. Another study on anomaly detection in fleet monitoring used unsupervised learning techniques, including K-means and hierarchical clustering, to identify abnormal behavior in time series data collected via the Fleet Management System (FMS) protocol. Parameters such as fuel consumption rate (liters per kilometer) were analyzed using data from 214 buses over six months. The results highlighted the algorithms' ability to detect outliers. However, challenges include incomplete data for some buses and a limited number of available signals, hindering more robust implementation [15].

Additionally, more advanced techniques based on Deep Learning have been explored, which utilize neural networks composed of multiple layers to extract complex patterns from large datasets. A predictive maintenance system was developed using an IoT architecture integrated with artificial intelligence, designed to monitor bus fleets. This system combines a multilayer perceptron artificial neural network (MLP-ANN) to predict engine wear with driver behavior analysis using the K-means clustering algorithm. The results indicate a low mean squared error (MSE), confirming the model's accuracy. The main limitations include the restricted capacity to connect additional sensors to the integrated hardware [16]. Another study compares remaining useful life (RUL) prediction models for turbocharger actuators in diesel engines, focusing on TabNet, RNN, and the Accelerated Weibull Failure Time (AWFT) model. AWFT demonstrated superior performance, achieving a concordance of 0.94 and approximately 15% lower error, making it suitable for defining fleet-specific preventive maintenance intervals. A combination of reliability data, electronic module records, and operational characteristics such as coolant temperature and average speed was used. Limitations included data dispersion and the difficulty of capturing random failures, which account for 10% of cases. Despite these challenges, the approach facilitates proactive replacements, reducing unexpected failures and enhancing system reliability [17].

As part of Industry 4.0, these techniques represent an innovative approach to overcoming the challenges of vehicle maintenance management. However, many of these models are considered "black boxes", meaning they produce predictions without offering a clear explanation of how they were reached. To overcome this limitation, Explainable Artificial Intelligence (XAI) has emerged as an evolving field with the primary goal of making AI models more transparent and understandable to humans, thereby enabling better insight into how and why certain decisions are made [18,19].

In the field of predictive maintenance, various studies have explored the use of XAI to improve the interpretability of their models, employing a range of techniques such as Shapley Additive Explanations (SHAP), Partial Dependence Plots (PDP), and Individual Conditional Expectation (ICE), among others [20–23]. Specifically for buses, the application of XAI in maintenance management has been examined in previous studies, albeit in a more limited manner. One study relied on data obtained from maintenance records and snapshots of the electronic control module during repair moments [24], while another merely presented the results of applying one of these techniques [25]. Additionally, both studies restricted their use to SHAP, without delving into other techniques that could provide different perspectives.

As noted, despite various studies focused on the application of Machine Learning and explainable artificial intelligence techniques in the field of maintenance management, their application in urban bus fleets using On-Board Diagnostics (OBDII) data remains virtually nonexistent. Furthermore, existing research tends to focus on a single technique, failing to explore a broader range of methods, leaving a significant gap in the literature. Given that modern buses are equipped with advanced technological infrastructure capable of generating and transmitting real-time data for subsequent processing and analysis, it is essential to present a comprehensive methodology that integrates Machine Learning with multiple XAI techniques. This approach would offer maintenance managers more robust tools to address fleet challenges, promoting more digitalized maintenance practices aligned with the principles of Industry 4.0.

To demonstrate the potential of these techniques in modern urban bus fleets using data from the OBDII, this study presents a use case focused on the regeneration process of the diesel particulate filter (DPF). The operation of the DPF under urban service conditions contributes to accelerated engine oil degradation, and the lack of documentation on the control system limits the maintenance department's ability to propose effective solutions. This study combines Machine Learning with XAI techniques to analyze the three states governing the DPF regeneration process, identifying the most influential variables from the electronic control unit (ECU) and highlighting the value of XAI in optimizing maintenance management. The article is structured as follows: Section 2 details the addressed problem and the available infrastructure; Section 3 describes the methodology employed and the tools used; Section 4 presents the results obtained and the proposed solutions; and finally, Section 5 discusses the study's conclusions.

2. Context of the Case Study

2.1. Accelerated Degradation in Engine Oil

Vehicles must comply with the European Emissions standard, which establishes, among other limitations, a limit on the amount of soot they can generate. The soot is captured in the diesel particulate filter (DPF), which is part of the exhaust after-treatment system. When the DPF is filled, the ECU requests a cleaning through the process of "regeneration" which involves increasing the exhaust gas temperature (from approximately 240 °C to 500 °C) to promote the burning of the particles trapped in the DPF (Figure 1). In service, a complete regeneration usually takes around 30 min, but the vehicle must be able to maintain such a temperature. To achieve this, high-demand operating conditions are required. If the vehicle cannot maintain these conditions, the regeneration will be stopped and postponed. These failed regeneration attempts negatively impact the engine oil, as the engine is forced to operate in suboptimal conditions multiple times.

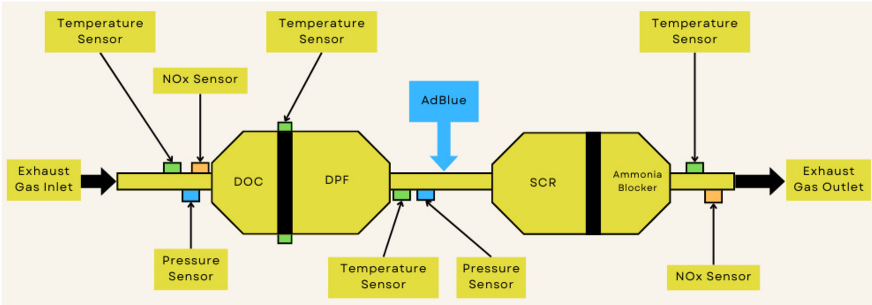

Figure 1. Overview of the After-treatment System Used in the Studied Bus.

For buses in urban service, it is difficult to achieve such conditions, which reduces the mileage between oil changes. Urban buses typically operate under demanding conditions characterized by frequent stops, low average speeds, and prolonged idling times. Consequently, maintenance—conducted in-house by the urban transport company—becomes substantially more costly, particularly given that this issue affects over 200 vehicles equipped with the after-treatment system.

To address the problem and propose solutions, the first step is to fully understand the states related to the regeneration process. From the ECU, it is possible to obtain the variable "Regeneration Status", a categorical variable that indicates the phase in which the process is. Expert knowledge provides additional information for most of the states (Table 1). However, detailed information on the operating conditions that trigger each state is lacking. Therefore, the objective is to analyze and comprehend these varying operating conditions to identify the specific triggers for each regeneration state.

Table 1. DPF regeneration status.

Status	Name	Information
2	Disabled	No regeneration
4	Temperature control	Change on engine operating conditions; only half of the cylinders are working alternately
8	Fuel injection	Seventh fuel injector enabled. Fuel is injected directly into the exhaust gases
32	Enabled	No information available

For this purpose, the study was conducted on a 12 m EURO VI-D diesel hybrid bus (Table 2) that primarily operates on intercity routes. This bus was selected due to its consistent operation on a single route, reducing variability and ensuring more reliable data analysis. Data collected over a 7-month period were used and obtained through the OBDII system.

Table 2. Engine Specifications of the EURO VI-D Diesel Hybrid Bus.

Maximum power [kW]@ 2200 rpm	220
Maximum torque [Nm]	1200
Bore Diameter [mm]	110
Displacement [cm^3]	7700
Stroke [mm]	135
Number of cylinders	6
Engine position	Vertical

2.2. Technological Infrastructure and Data Collection

The ECU of the after-treatment system can provide 99 variables to explore, related not only to the DPF but also to other components of the system, such as:

- Diesel oxidation catalyst (DOC): Converts pollutants like carbon monoxide and hydrocarbons into carbon dioxide and water.
- Selective catalytic reduction (SCR): A system that uses urea to transform nitrogen oxides (NOx) into nitrogen and water.
- Ammonia blocker: A filter designed to prevent the release of residual ammonia in the SCR system, ensuring that only clean gases are emitted into the environment.

These variables include parameters such as temperatures and pressures in the exhaust gases and NOx measurements, among others. Additionally, there are calculated variables by electronics, such as the time since the last regeneration.

After reviewing the literature on the phenomenon and historical alarms presented in the ECU, and with the support of expert knowledge, the relevant variables that could provide insights into the different states of the process were selected. Initially, 12 variables from the after-treatment system's ECU were selected for this purpose (Table 3).

Table 3. Selected variables for the case study.

Variables	Unit
SCR NOx outlet	ppm
SCR NOx inlet	ppm
DPF Backpressure inlet	mbar
DPF Backpressure outlet	mbar
DPF regen state	-
Engine speed	rpm
DOC inlet temp	°C
DPF outlet temp	°C
SCR outlet temp	°C
DOC outlet temp	°C
Speed	km/h
Duration	min

The available telematics and storage infrastructure, provided by a third-party company, facilitated the downloading of data in CSV format for subsequent processing and analysis. The data were analyzed using the Python programming language along with libraries such as pandas, matplotlib, and scikit-learn. One of the most significant limitations encountered during data analysis was the sampling frequency, which is set at 3 min.

3. Materials and Methods

In the field of Machine Learning, there are different approaches. In this study, the most suitable methodology was selected based on the nature of the available data. For our case, we have a dataset of variables (known as features) that includes operational parameters extracted from the OBDII system, such as engine speed and vehicle speed, among others. Additionally, we have an output variable or target, which, in this case, represents the different regeneration states in the DPF. Given this context, a supervised learning approach focused on classification was chosen, as it allows modeling the relationship between the input variables and the system states during the regeneration process. Alternative approaches, such as unsupervised learning and regression, were also considered but found unsuitable for this study. Unsupervised learning methods, while useful for exploratory analysis and clustering, do not leverage the labeled target variable available in this dataset, limiting their applicability in this context. Similarly, regression-based methods were not appropriate because the target variable is categorical and represents nominal states of the regeneration process, which are best addressed through classification techniques. The choice of supervised classification ensures that the methodology aligns with the nature of the data and the study's objectives.

3.1. One vs. Rest

As previously mentioned, the study problem involves supervised classification, meaning the model is trained on a labeled dataset where each instance is accompanied by a known label or class. This implies that for specific values of variables such as speed and engine regime, there is a particular regeneration process state associated. Additionally, the target label is a nominal multiclass variable, which means the target variable can take on different values, representing each state of the DPF regeneration process. These values do not follow a hierarchical order; rather, each class is a unique and independent category. An example of a nominal classification would be a vehicle's fuel type (gasoline, diesel, or hydrogen), where the categories have no intrinsic order.

Given that the objective is to understand the specific conditions that promote the activation of each state, a One-vs-Rest (OvR) approach was chosen. This strategy transforms a multiclass problem into several binary classification problems, allowing each class to be individually compared to all others. In this case study, this involves generating three binary models, each focused on distinguishing a specific DPF regeneration state from the other possible states:

- Model 1: Classifies Regeneration State 4 against all other states.
- Model 2: Classifies Regeneration State 8 against all other states.
- Model 3: Classifies Regeneration State 32 against all other states.

The One-vs-Rest approach was selected to provide a comprehensive analysis of the conditions that activate each specific state of the DPF regeneration process. As previously mentioned, the OvR method breaks the problem into simpler binary classification tasks, allowing each state to be analyzed independently against all others. This decomposition offers several advantages: it improves interpretability by isolating the unique features and operating conditions that influence the activation of each state, facilitates the identification of patterns that might be obscured in a global multiclass approach, and enables customized tuning for each binary classifier.

3.2. Exploration Data Analysis and Preprocessing

Once the target variable was defined, the first step involved an exploratory data analysis (EDA), a critical phase in Machine Learning projects that often accounts for a significant portion of the overall project timeline. This process helps identify and understand the most relevant data characteristics and makes necessary adjustments to enhance the effectiveness of the models being developed.

Initially, data points not directly related to the regeneration process, such as those where regeneration was inactive (State 2), were discarded. This decision allows the analysis to focus on moments of interest—specifically, those when regeneration occurs—and additionally helps avoid diluting relevant patterns with data from normal operating conditions. Moreover, certain data points deemed anomalous, particularly those representing regeneration events of unusually short durations, were removed, as they might indicate errors in data collection or events that do not reflect regular DPF operation.

A feature engineering process was also carried out, which is a critical step in Machine Learning that involves transforming raw data into meaningful features to enhance a model's predictive capacity. By leveraging domain expertise and empirical analysis, this process aims to enrich the dataset by creating variables that effectively represent key system dynamics, capturing underlying patterns and relationships that might otherwise remain hidden. For instance, it facilitates the integration of temporal patterns or derived metrics, adding valuable context to the system under analysis. While feature engineering often results in significant performance improvements—such as increased model accuracy and interpretability—it also presents challenges. Specifically, creating excessively derived features can increase model complexity, thereby heightening the risk of overfitting and reducing generalizability.

Finally, visualizations such as heatmaps were utilized to analyze potentially strong correlations between variables using Pearson's correlation coefficient. This tool was particularly useful for evaluating the variables related to temperature in the after-treatment system, where high correlations were expected due to the proximity and arrangement of the sensors. This issue is relevant because strong interdependence between variables can impact the model's performance by introducing redundancy and potentially causing overfitting. For this reason, the variable with the strongest correlation to the target was selected, ensuring that the model worked with the most relevant features while reducing multicollinearity.

3.3. Classification Algorithms

To address the supervised classification problem in this study, several Machine Learning algorithms were implemented. These algorithms, which have been previously demonstrated as effective tools for classification tasks in predictive maintenance [4,26–28], were evaluated using different hyperparameter combinations through grid search, selecting the most optimal configuration for each model. The algorithms tested included XGBoost (version 2.1.2), LightGBM (version 4.5.0), Random Forest, and Support Vector Machine (Scikit learn version 1.5.1). For each state, the algorithm demonstrating the best performance was selected.

- XGBoost: Extreme Gradient Boosting is an advanced Machine Learning algorithm that builds on decision tree models to achieve high predictive accuracy. It is known for its efficiency and scalability, making it suitable for handling large datasets. XGBoost introduces unique features like sparsity awareness, which improves performance on datasets with missing values, and a weighted quantile sketch, allowing it to efficiently manage weighted data. Its optimized design enables fast computations and parallel processing, making it a popular choice for data scientists tackling complex classification and regression tasks. For XGBoost, the key hyperparameters adjusted during the tuning process were n_estimators, max_depth, and learning_rate.
- LightGBM: Light Gradient-Boosting Machine is a highly efficient Machine Learning algorithm designed to handle large-scale data by using gradient boosting with decision trees. Developed by Microsoft, it introduces methods like Gradient-based One-Side Sampling (GOSS) and Exclusive Feature Bundling (EFB) to speed up training and reduce computational demands without sacrificing accuracy. These techniques enable LightGBM to work efficiently with large datasets and high-dimensional feature spaces, making it a popular choice for tasks like classification. In the case of LightGBM, the hyperparameters optimized included n_estimators, max_depth, and learning_rate.
- Random Forest: The Random Forest algorithm is a widely used ensemble method that improves accuracy and robustness by combining the predictions of multiple decision trees. Each tree is built using a randomly selected subset of the data and features, reducing the risk of overfitting and enhancing the model's resilience to

noisy data. This method makes Random Forests particularly effective for handling datasets with a high number of input variables, making them suitable for complex classification and regression tasks. For Random Forest, the hyperparameters tested included n_estimators, max_depth, and min_samples_split.

- SVM: Support Vector Machines are a classification method used to separate data into distinct classes by finding an optimal boundary. This boundary is defined as a hyperplane that maximizes the margin between different classes, allowing for better separation and classification accuracy. SVMs can also handle complex, non-linear data by mapping inputs into a higher-dimensional space through kernel functions, enabling them to find the best-separating hyperplane even when classes overlap in lower-dimensional spaces. For SVM, the primary hyperparameters tuned were C and kernel.

To evaluate the performance of the classification algorithms, metrics such as confusion matrix, precision, recall, and F1-score were employed, providing insights into the effectiveness of the models' predictions. Precision measures the proportion of correctly predicted positive instances out of all predicted positives, while recall (or sensitivity) reflects the model's ability to correctly identify actual positive instances. The F1-score, a balanced metric combining precision and recall, was used as the primary criterion to determine the best-performing algorithm in each case.

$$Precision = \frac{True\ Positive\ (TP)}{True\ Positive\ (TP) + False\ Positive\ (FP)}$$

$$Recall = \frac{True\ Positive\ (TP)}{True\ Positive\ (TP) + False\ Negative\ (FN)}$$

$$F1\ Score = 2 \times \frac{Precision\ \times\ Recall}{Precision + Recall}$$

3.4. Explainability Techniques

Starting with SHAP (Shapley Additive Explanations), this explainability technique is based on cooperative game theory and provides a robust framework for understanding the contribution of each feature to the outcome of a prediction. This technique is highly versatile and applicable to a wide range of models, including tree-based models, linear regressions, and deep neural networks. SHAP enables the generation of local explanations by calculating SHAP values for individual data instances. These values precisely quantify how much each feature contributes to the outcome of a specific prediction. Furthermore, by aggregating SHAP values from multiple instances, global explanations are obtained, offering a comprehensive view of the model's behavior while highlighting the relative importance and average impact of each feature across the entire dataset [20,29,30]. Applied to the studied problem, SHAP analyzes the operational data of the bus, such as vehicle speed, engine regime, and pressure at the DPF inlet, among other key variables. Each of these features is evaluated to determine its influence on the model's prediction. This influence is quantified through SHAP values, which decompose the prediction into contributions attributable to each feature. For instance, SHAP can identify that, for the activation of a specific state, the DPF inlet pressure and engine regime are critical factors. Additionally, SHAP not only measures the direct impact of each feature but also considers the effect of its absence, enabling the evaluation of its relative importance in the context of the other variables.

Other relevant techniques include Partial Dependence Plots (PDP) and Individual Conditional Expectation (ICE) plots, which are data visualization and statistical analysis tools used to understand how a specific feature affects the predicted outcome of a model while keeping the other features constant. PDPs show the average effect of a feature on the target variable, allowing for a general analysis of the relationship between the feature and the model's predictions. In contrast, ICE plots provide a more detailed and granular view by displaying multiple curves, each representing the feature's impact on the outcome for individual data instances. ICE plots enable the observation of individual variations and potential non-linear effects or interactions, while PDPs summarize the general behavior of the model concerning a particular feature [20,29,30]. For the case study, these techniques evaluate a specific parameter, such as bus speed, by varying it while keeping other variables like engine regime and temperature constant. PDPs would then show how changes in speed affect, on average, the probability of activating the specific state in that model, while ICE plots would allow for the analysis of whether these relationships are consistent across different observations or significantly vary depending on the context.

4. Results

4.1. Exploration Data Analysis

As previously mentioned, the first phase of the analysis focused on data cleaning to ensure that the information used was both relevant and representative of the phenomenon under study. To achieve this, data points corresponding to State 2 of the regeneration process, which is associated with the inactive condition, were removed (Table 4). This state does not provide significant information about the active conditions of the process and, therefore, is not relevant to the objective of the study.

Table 4. Data Distribution by Percentage for Each Regeneration State.

DPF Regeneration State	DPF State Data Distribution (%)
2	96.49
8	2.56
4	0.79
32	0.17

Subsequently, anomalous data points (outliers) were identified and removed. In this case, outliers were defined as records where the regeneration process was activated and deactivated within a 3-min interval. This behavior does not reflect the normal functioning of the system, which typically takes approximately 20 to 40 min, and could result from data collection errors or atypical, non-representative conditions. This cleaning process led to the removal of 16 points, ensuring that subsequent analyses focused solely on normal operating conditions.

The next step involved analyzing correlations among key system variables. The after-treatment system includes multiple temperature sensors distributed throughout the mechanism. Due to the physical proximity of these sensors and the thermal nature of the phenomenon, high correlations among these variables were expected. This was confirmed in Figure 2, which shows a considerable correlation between the various recorded temperature variables.

In Machine Learning, high correlation between variables can lead to redundancies and negatively impact the stability of predictive models. For this reason, only one temperature variable was selected for inclusion in the study. The selection criterion was the correlation with the target variable, which, in this case, represents the DPF regeneration state. This target variable had already been transformed into three binary variables using the one-vs-rest approach, representing the different regeneration states.

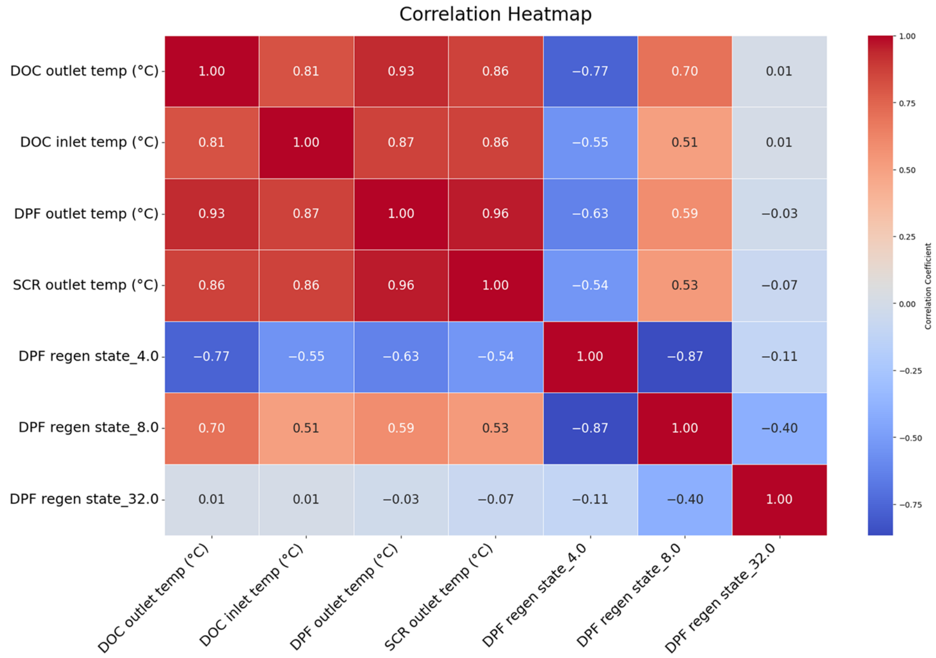

Figure 2. Correlation analysis for the temperatures in the after-treatment system.

After analysis, it was determined that the temperature recorded at the outlet of the diesel oxidation catalyst (DOC) exhibited the highest correlation with the regeneration states (Figure 3). Therefore, this variable was included in the final dataset for subsequent analyses.

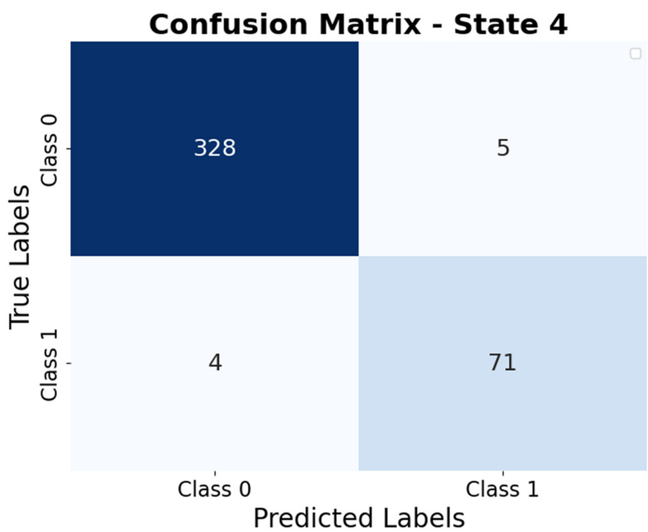

Figure 3. Confusion Matrix for State 4.

The final step involved feature engineering, which consists of creating new variables to better capture the relationships between existing variables and their influence on the regeneration process. Specifically, it was considered that temperature increases might have a significant impact on system behavior and the activation of regeneration states. To reflect this phenomenon, a new variable was created to measure the rate of temperature change over time for each recorded instance. This variable was calculated as follows:

$$Temp\ Change\ Rate = \frac{\Delta TempDOC}{\Delta Duration}$$

where $\Delta TempDOC$ represents the temperature difference between two consecutive instances and $\Delta Duration$ is the corresponding time interval. This new feature allows for a more precise analysis of the system's thermal dynamics and their influence on the regeneration states.

Additional variables were calculated to better capture the system's dynamics. These include the pressure delta in the DPF and the NOx delta in the SCR. The pressure delta in the DPF provides critical information about the resistance to gas flow through the filter, serving as a key indicator for estimating DPF clogging. Meanwhile, the NOx delta in the SCR assesses the efficiency of the selective catalytic reduction system, offering a precise measure of the change in NOx concentration before and after the catalyst.

4.2. Model Development

The results from applying the aforementioned algorithms are presented in Figures 3–5 and Table 5, with the optimal hyperparameters detailed in Table 6. As observed, the models generally perform better when predicting classes where State 8 has greater representation, directly influenced by the data imbalance. For this reason, the F1-score for the minority class was prioritized as the key metric, as it provides a more precise evaluation of the model's performance in underrepresented cases. The specific results for each state are as follows:

- State 4: The LightGBM model demonstrated slightly better performance than XGBoost in terms of F1-score for predicting the minority class (Class 1).
- State 8: For this state, the LightGBM model also demonstrated the best results in terms of F1-score for the minority class (class 0).
- State 32: For this state, the metrics for the minority class (Class 1) were significantly impacted by the data imbalance. However, the results obtained with the LightGBM model, which achieved the best metrics among the evaluated models, remain valuable. These results can help identify preliminary patterns in the regeneration process associated with this state through XAI.

Table 5. Performance Metrics for Each Model across DPF Status.

State	Model Name	Class 0 Precision	Class 0 Recall	Class 0 F1-Score	Class 1 Precision	Class 1 Recall	Class 1 F1-Score
4	Random Forest	0.9703	0.9820	0.9761	0.9155	0.8667	0.8904
	LightGBM	0.9850	0.9850	0.9850	0.9333	0.9333	0.9333
	XGBoost	0.9848	0.9730	0.9789	0.8861	0.9333	0.9091
	SVM	0.9788	0.9700	0.9744	0.8718	0.9067	0.8889
8	Random Forest	0.9259	0.7979	0.8571	0.9419	0.9809	0.9610
	LightGBM	0.9130	0.8936	0.9032	0.9684	0.9745	0.9714
	XGBoost	0.9419	0.8617	0.9000	0.9596	0.9841	0.9717
	SVM	0.7596	0.8404	0.7980	0.9507	0.9204	0.9353
32	Random Forest	0.9792	0.9666	0.9728	0.4583	0.5789	0.5116
	LightGBM	0.9818	0.9717	0.9767	0.5217	0.6316	0.5714
	XGBoost	0.9794	0.9769	0.9781	0.5500	0.5789	0.5641
	SVM	0.9831	0.8946	0.9367	0.2407	0.6842	0.3562

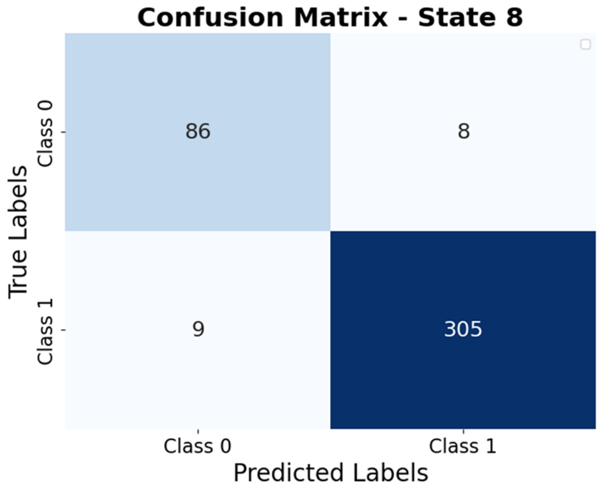

Figure 4. Confusion Matrix for State 8.

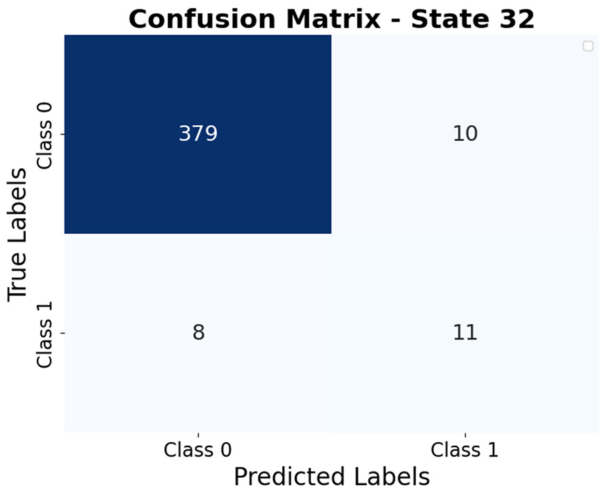

Figure 5. Confusion Matrix for State 32.

Table 6. Hyperparameter Tuning Results for Each Model and Status.

State	Model Name	Best Params
4	Random Forest	{'max_depth': 10, 'min_samples_split': 4, 'n_estimators': 140}
	LightGBM	{'learning_rate': 0.1, 'max_depth': 10, 'n_estimators': 120}
	XGBoost	{'learning_rate': 0.1, 'max_depth': 6, 'n_estimators': 130, 'scale_pos_weight': 10}
	SVM	{'C': 10, 'kernel': 'rbf'}
8	Random Forest	{'max_depth': 10, 'min_samples_split': 2, 'n_estimators': 120}
	LightGBM	{'learning_rate': 0.1, 'max_depth': 12, 'n_estimators': 110}
	XGBoost	{'learning_rate': 0.1, 'max_depth': 7, 'n_estimators': 120, 'scale_pos_weight': 10}
	SVM	{'C': 1, 'kernel': 'rbf'}
32	Random Forest	{'max_depth': 10, 'min_samples_split': 4, 'n_estimators': 120}
	LightGBM	{'learning_rate': 0.01, 'max_depth': 10, 'n_estimators': 150}
	XGBoost	{'learning_rate': 0.1, 'max_depth': 6, 'n_estimators': 100, 'scale_pos_weight': 10}
	SVM	{'C': 1, 'kernel': 'rbf'}

4.3. Results from XAI Techniques

4.3.1. SHAP

SHAP offers a versatile range of tools to decompose and quantify the impact of each variable on the model's predictions, thus facilitating the understanding of complex decisions. In this study, two key representations have been utilized:

- Beeswarm Plot: This tool visualizes the impact of all variables on the predictions in an aggregated manner. Each point in the plot represents an instance, with colors reflecting the magnitude of the variables (ranging from low to high). Additionally, variables are sorted from the highest to lowest influence on the model, making it easier to identify which factors have the most significant impact on predictions.
- Dependence Plot: This SHAP tool provides an in-depth analysis of how a specific variable affects the model's predictions, also considering its interaction with other variables. To enhance the interpretability of the presented plots, the scaling applied during model generation was reversed. However, it is essential to note that the SHAP values were generated using scaled data.

Figures 6 and 7 show the Beeswarm Plots for the models predicting States 4 and 8. In both cases, the temperature at the DOC emerges as the most influential factor in predicting these states. For State 8, high-temperature values are associated with its activation, whereas for State 4, the activation occurs at lower temperature values.

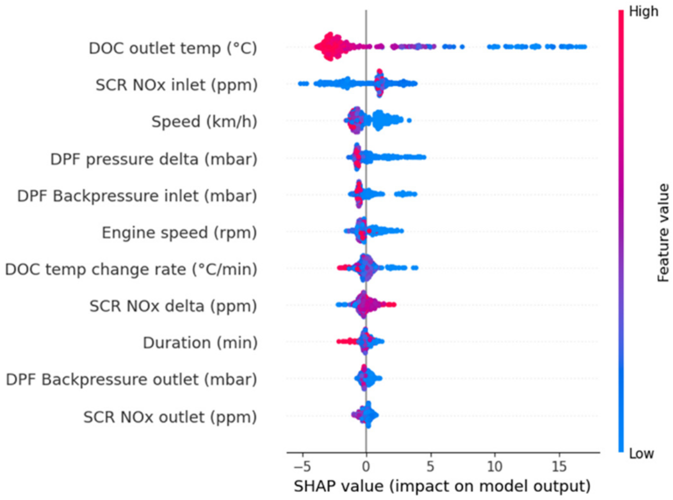

Figure 6. Feature Contribution Analysis Using Beeswarm Plot (State 4).

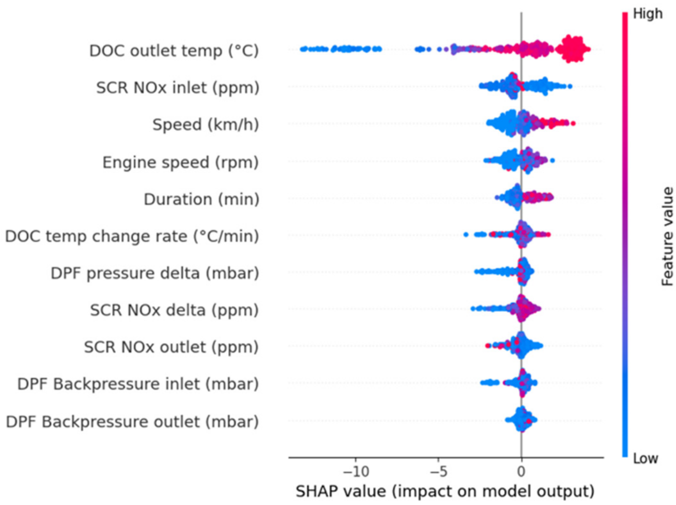

Figure 7. Feature Contribution Analysis Using Beeswarm Plot (State 8).

For State 4, low values of speed and DPF pressure delta significantly contribute to the activation of this state. Conversely, State 8 displays the opposite behavior, with higher values of speed and engine regime associated with its activation. Another variable with a defined behavior for this state is duration, where higher values tend to favor its activation.

NOx in the SCR inlet is the second most influential variable for States 4 and 8. However, its behavior is peculiar, as it seems to exhibit low values throughout the range, with a cluster of high points. To examine this behavior in more detail, Figures 8 and 9 present the SHAP dependence plots for this variable, where the interaction with the DPF pressure delta is also highlighted. The cluster of high points is associated with elevated DPF pressure delta values, which makes physical sense considering that a higher-pressure delta implies a greater gas flow and, consequently, a higher NOx reading. A significant insight is that during these high-demand moments, likely linked to acceleration events, the influence of the variable (as represented by SHAP values) tends to approach zero. This suggests that under such conditions, the model no longer considers NOx at the SCR inlet as a key factor in predicting the states.

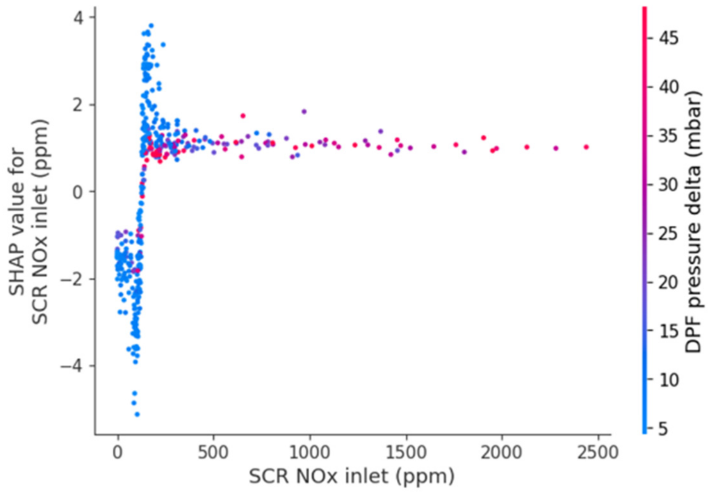

Figure 8. SHAP Dependence Plot for NOx inlet in the SCR with DPF Pressure Delta (State 4).

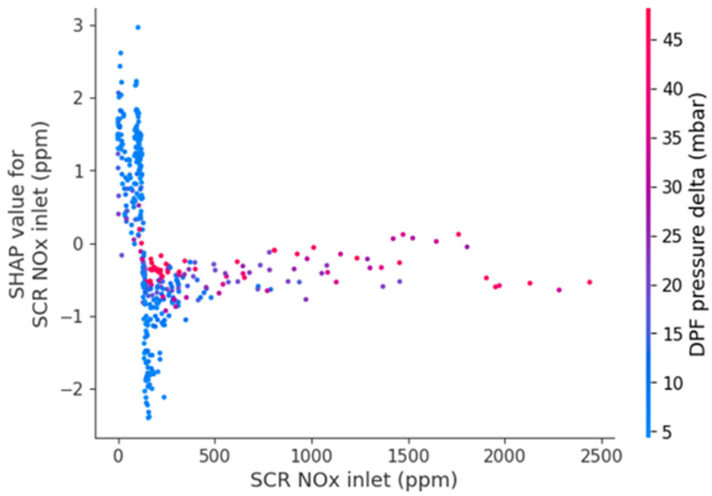

Figure 9. SHAP Dependence Plot for NOx inlet in the SCR with DPF Pressure Delta (State 8).

Regarding State 32, the DOC outlet temperature also emerged as the most influential variable (Figure 10). However, its impact on the prediction is less straightforward compared to States 4 and 8. In this case, both the activation and non-activation of the state are associated with high and low-temperature values, reflecting a more complex behavior. Additionally, variables related to duration and the NOx delta also have a significant influence on the prediction. Higher values of both variables are predominantly associated with the non-activation of State 32.

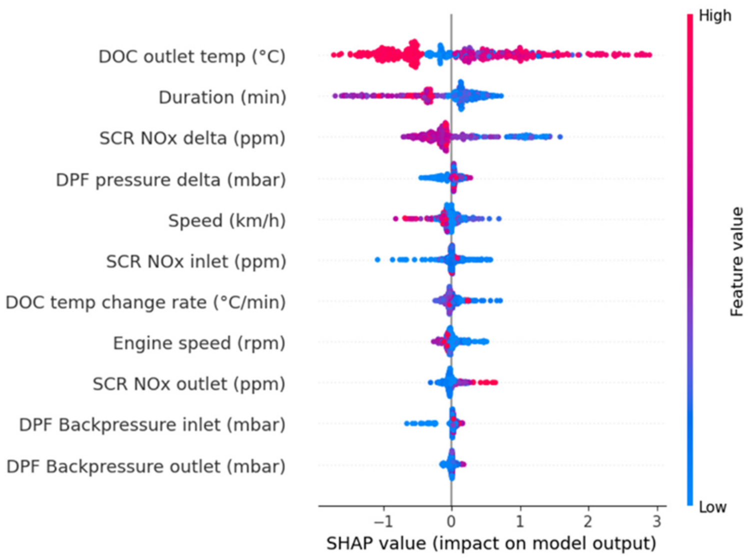

Figure 10. Feature Contribution Analysis Using Beeswarm Plot (State 32).

The dependence plot for the NOx delta (Figure 11) reveals interesting behavior. Multiple data points exhibit negative NOx delta values, which lack physical plausibility, as this metric should represent a decrease or, at worst, remain constant. However, these negative values appear to be associated with the activation of State 32, showing a clear trend: more pronounced negative values are linked to positive SHAP values. Additionally, the analysis incorporates duration as a secondary variable (color axis), corroborating the findings from the Beeswarm plot: low duration values are associated with the activation of this state. By combining these two variables, it becomes evident that State 32 is influenced by these initial moments where significantly negative NOx delta values are recorded.

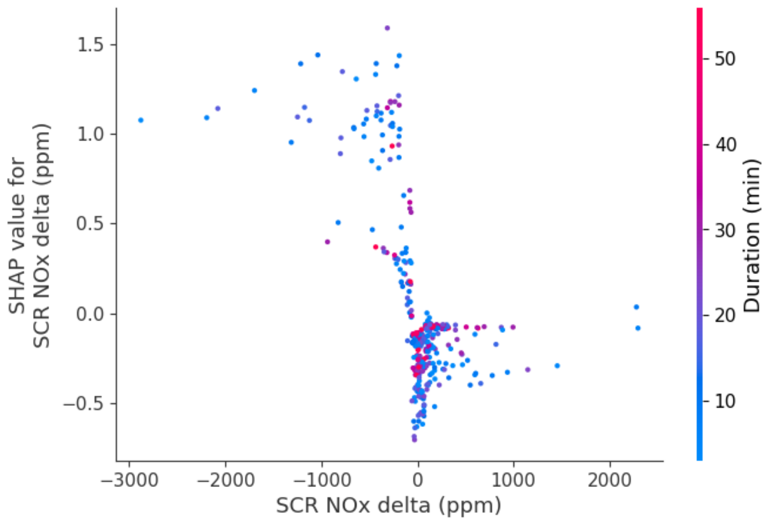

Figure 11. SHAP Dependence Plot for NOx Delta in the SCR with Duration (State 32).

Regarding DOC temperature, SHAP values show two notable peaks: one around 350 °C and another above 500 °C, followed by a cluster of points driving the prediction towards non-activation within the 550–600 °C range (Figure 12). These activation peaks appear to be linked to the previously observed low NOx delta values. Despite the analysis conducted using several dependence plots, this behavior in temperature could not be conclusively explained. It is important to consider that, as mentioned earlier, the model's performance shows room for improvement, with metrics that have yet to reach optimal levels. This limited performance may hinder the model's ability to identify clear and consistent patterns, resulting in predictions that are not entirely interpretable. Consequently, these atypical behaviors could largely stem from the lack of representativeness in the training data for the minority class (activation of this state).

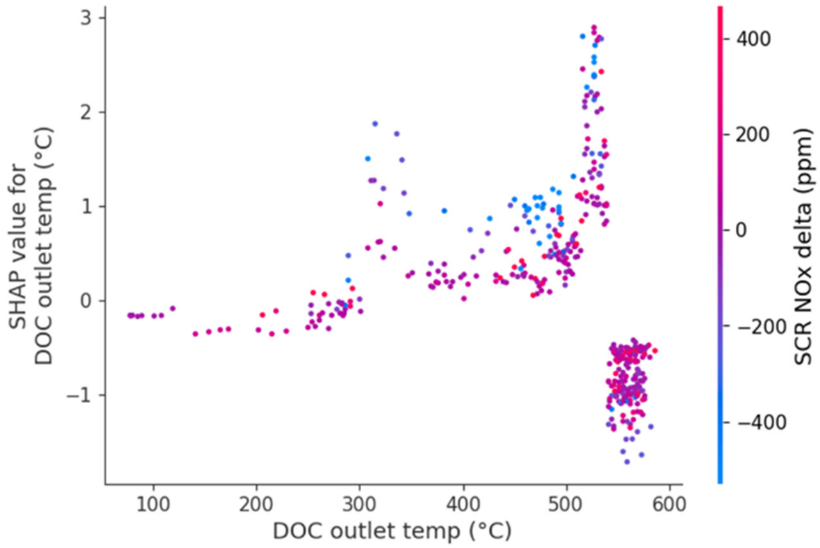

Figure 12. SHAP Dependence Plot for Outlet Temperature in the DOC with NOx Delta in the SCR (State 32).

Given that State 8 is clearly associated with high temperatures in the system—conditions that are ideal for regeneration—the goal is to identify controllable variables that can promote its activation. These variables include vehicle speed, engine regime, and torque, the latter being indirectly represented by the pressure delta in the DPF. From the Beeswarm analysis, it has been shown that both speed and regime have a considerable impact on state activation.

Figures 13 and 14 present the Partial Dependence Plots for these variables. For speed, it is observed that its influence remains minimal up to 30 km/h; however, beyond this threshold, its impact becomes significant, promoting state activation. Regarding the engine regime, parabolic behavior is detected after idling values (~600 rpm). There is a decrease in influence around 800 rpm, followed by a notable increase in impact between 1000 and 1200 rpm.

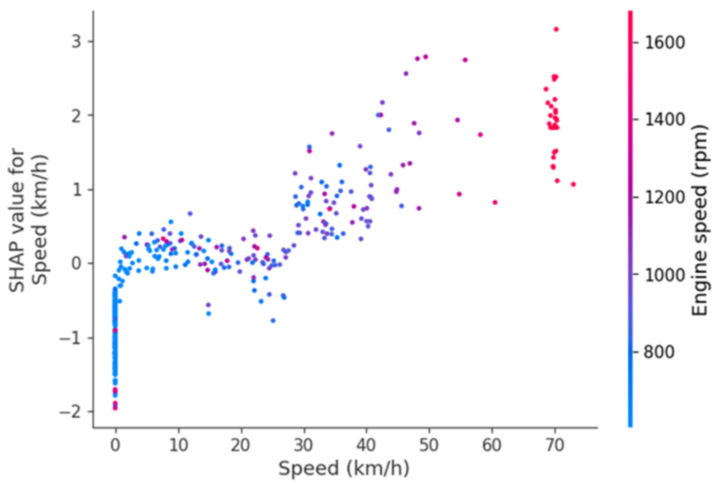

Figure 13. SHAP Dependence Plot for Speed with Engine Speed (State 8).

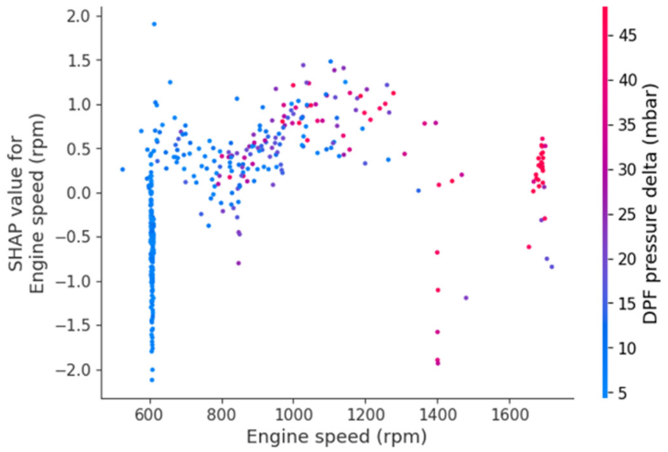

Figure 14. SHAP Dependence Plot for Engine Speed with DPF Pressure Delta (State 8).

4.3.2. PDP/ICE

Figures 15–18 present the univariate plots obtained by applying Partial Dependence Plot (PDP), represented by the orange line (average), and Individual Conditional Expectation (ICE), represented by the blue lines, for DOC temperature and DPF pressure delta for States 4 and 8. Previous analyses using SHAP identified DOC temperature as a key factor in predicting these states. This behavior is reaffirmed in the presented plots, where the partial dependence plot of DOC temperature ranges from values close to 1 to approximately 0, being the only variable that displays such a wide range in terms of average impact. It is important to note that the temperature values correspond to the scaled data used during model generation. Regarding the pressure delta and the rest of the variables, a common behavior is observed, stemming from the inherent limitation of these techniques (results are obtained by modifying only the variable of interest while keeping the rest constant). PDP assumes a premise of independence between variables, which does not always hold true for the analyzed data. This limitation is evident in the ICE lines, which exhibit varied behaviors and divergent trajectories depending on individual instances. When these trajectories are averaged to generate the PDP, the variable's impact may appear diminished or smoothed, making it more challenging to assess its isolated influence.

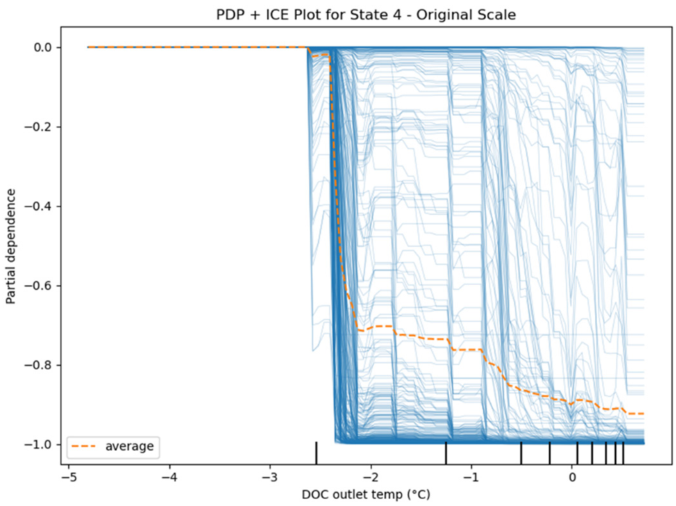

Figure 15. PDP and ICE Plot for Outlet Temperature in the DOC (State 4).

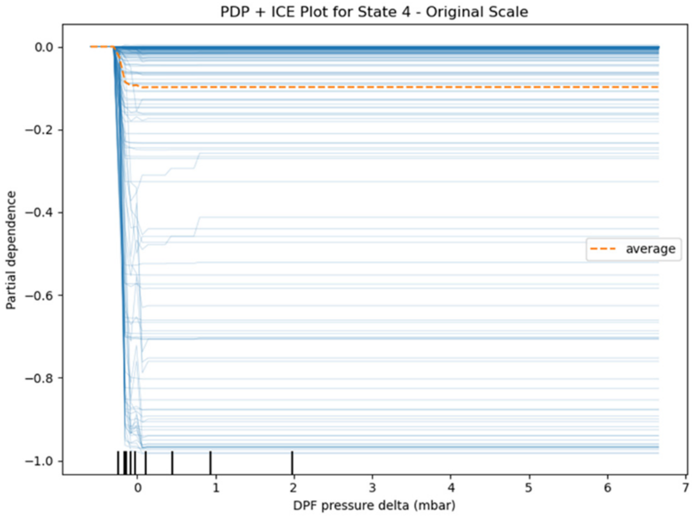

Figure 16. PDP and ICE Plot for DPF Pressure Delta (State 8).

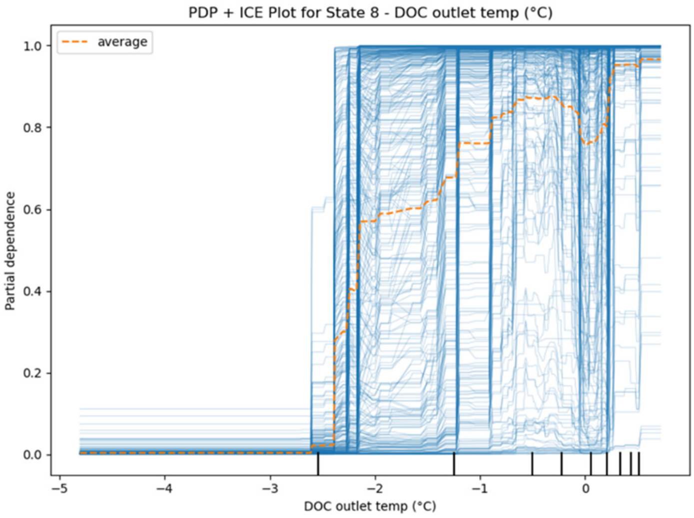

Figure 17. PDP and ICE Plot for Outlet Temperature in the DOC (State 8).

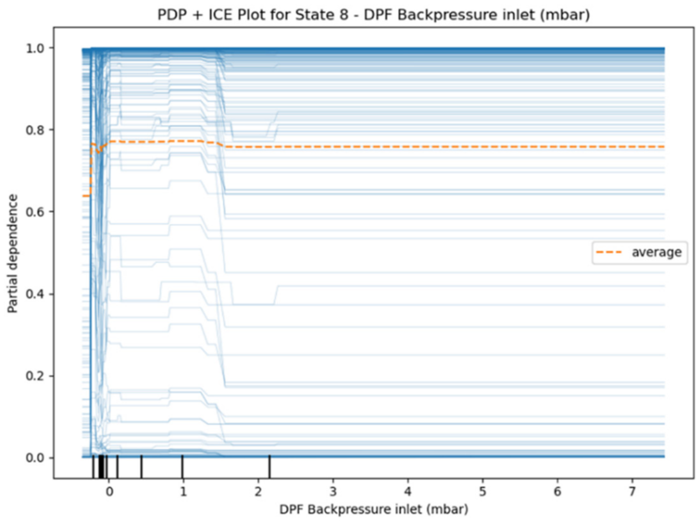

Figure 18. PDP and ICE Plot for DPF Backpressure inlet (State 8).

Figures 19 and 20 present the bivariate plots of this technique, which allow for the analysis of interactions between the two variables. This approach partially overcomes one of the previously mentioned limitations related to the lack of independence among variables. The presented plots show the combination of DOC temperature with vehicle speed for States 4 and 8. In both cases, a region is identified where speed does not appear to significantly influence the probability of state activation. This finding aligns with previous analyses conducted using SHAP and univariate PDP, which highlighted temperature as the most determining factor in predicting these states. However, in the intermediate zones of the plots, distinct patterns in the form of steps begin to emerge, confirming that speed plays a more significant role in activation. Table 7 serves as a reference for interpreting the scale of the variables used in the plots. Specifically, in the plot corresponding to State 8, it is identified that for speeds around 26 km/h, the impact of this variable starts to have a greater influence on the prediction of the state. This effect becomes more pronounced as speed reaches values close to 40 km/h, where its influence on the prediction becomes even more significant.

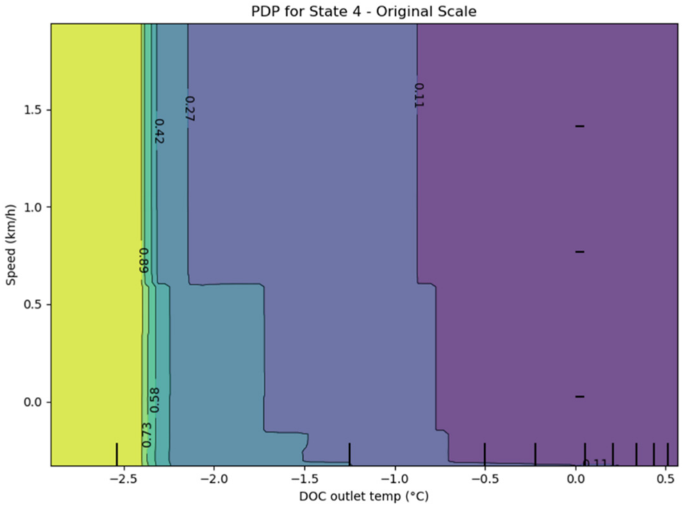

Figure 19. Bivariate PDP for Speed and Outlet Temperature in the DOC (State 4).

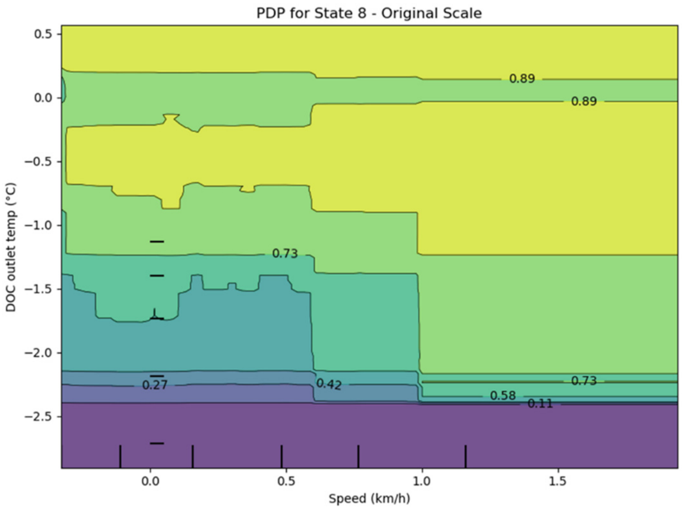

Figure 20. Bivariate PDP for Speed and Outlet Temperature in the DOC (State 8).

Table 7. Original and Scaled Values for DOC Outlet Temperature and Speed.

Percentile	DOC Outlet Temp (°C) Scaled	DOC Outlet Temp (°C) Original	Speed (km/h) Scaled	Speed (km/h) Original
10%	−2.5207	288	−0.3268	0
15%	−1.9217	343	−0.3268	0
20%	−1.2413	406	−0.3268	0
25%	−0.6875	457	−0.3268	0
30%	−0.4989	474	−0.3268	0
35%	−0.3212	490	−0.3268	0
40%	−0.2196	500	−0.3268	0
45%	−0.0745	513	−0.2615	2
50%	0.0543	525	−0.1095	7
55%	0.1087	530	0.0062	10
60%	0.2065	539	0.156	15
65%	0.2886	547	0.3748	22
70%	0.3370	551	0.4818	25
75%	0.3804	555	0.6443	30
80%	0.4348	560	0.7662	34
85%	0.4886	565	0.9475	39
90%	0.5109	567	1.1576	46

4.4. Improvement Proposal

Given that State 8 can be facilitated by controllable variables such as vehicle speed and engine regime, a strategic improvement plan is proposed to optimize the conditions for regeneration within the system:

- Rotation between service routes: For urban transport operators with routes that have varying average speeds, it is recommended to implement a strategic rotation of buses across these routes. This would ensure that vehicles regularly have the opportunity to achieve the necessary conditions for proper regeneration processes with greater frequency.
- Adjustment of the gearbox configuration: Vehicle gearboxes come with different operating configurations, which dictate the revolutions at which the vehicle shifts up or down. Currently, the buses operate in an economy mode, causing gear changes at relatively low revolutions (around 900 rpm). To promote higher demand for the vehicle, the proposal involves modifying the gear shift strategy in the buses, allowing them to reach higher revolutions and, consequently, the ideal range to facilitate regeneration.

Both approaches have the potential to significantly enhance regeneration conditions, although they present certain challenges that must be addressed prior to implementation. In the case of route rotation, it is critical that cities have a sufficient number of routes with high average speeds to effectively alternate buses. On the other hand, with modifying gearbox configuration, a detailed analysis is needed to assess whether the reduction in engine oil degradation outweighs the additional fuel consumption that this strategy might entail.

5. Conclusions and Future Work

Maintenance management in urban bus fleets has significant potential to benefit from the integration of advanced data analytics and explainable artificial intelligence (XAI). By leveraging Machine Learning and interpretability techniques, this study provided a framework for extracting valuable insights from OBDII data, enabling the identification of key factors that influence the activation of critical states in diesel particulate filter (DPF) regeneration processes. The results emphasize the dominant role of DOC outlet temperature in predicting DPF States 4 and 8, as demonstrated through SHAP, PDP, and ICE analyses. Notably, State 8 was found to strongly correlate with high temperatures—ideal conditions for effective regeneration. Additionally, the study identified actionable thresholds for operational variables, such as vehicle speed (e.g., above 30 km/h) and engine regime (e.g., 1000–1200 rpm), which align with the conditions favorable for State 8 activations. Two key strategies were proposed to improve regeneration outcomes: (1) the strategic rotation of buses between routes with varying speed profiles, which would only require coordinating the required distribution with the responsible department, and (2) modifying the operating mode of the automatic gearbox to enable the vehicle to reach higher engine speeds before shifting to the next gear. However, these strategies are not without challenges. Route rotation requires an adequate distribution of high-speed routes, while changes to the gearbox configuration demand careful consideration of fuel consumption impacts versus potential reductions in engine oil degradation.

This study highlights the value of integrating XAI techniques with operational data to derive actionable insights, even in the presence of data imbalances—a common challenge in data-driven approaches. Considering that the tools used to apply these techniques are open source, their applicability is primarily constrained by the requirement for specialized technical expertise in data analysis. Future efforts should prioritize real-world validation of these findings and focus on expanding the dataset to better explore State 32, where the current results indicate room for improvement. By adopting these tools and strategies, maintenance management of urban bus fleets can advance toward more efficient, reliable, and sustainable maintenance practices.

Author Contributions: Conceptualization, B.T. and J.L.C.; methodology, R.S.-M. and B.P.; software, R.S.-M.; validation, R.S.-M., B.T. and J.L.C.; formal analysis, R.S.-M. and B.P.; investigation, R.S.-M.; resources, B.T.; data curation, R.S.-M.; writing—original draft preparation, R.S.-M.; writing—review, and editing, B.T. and B.P.; visualization, R.S.-M.; supervision, B.T.; project administration, B.T.; funding acquisition, B.T. All authors have read and agreed to the published version of the manuscript.

Funding: This work has been partially supported by grant PID2020-119691RB-100 funded by MCIN/AEI/10.13039/501100011033.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in the study are openly available in Mendeley Data at DOI: 10.17632/jcf828mv7m.1.

Acknowledgments: We would like to thank the EMT, ALSA, and Jaltest staff for their support during the elaboration of this research.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Yan, J.; Meng, Y.; Lu, L.; Li, L. Industrial Big Data in an Industry 4.0 Environment: Challenges, Schemes, and Applications for Predictive Maintenance. *IEEE Access* **2017**, *5*, 23484–23491. [CrossRef]
2. Teoh, Y.K.; Gill, S.S.; Parlikad, A.K. IoT and Fog-Computing-Based Predictive Maintenance Model for Effective Asset Management in Industry 4.0 Using Machine Learning. *IEEE Internet Things J.* **2023**, *10*, 2087–2094. [CrossRef]
3. Resende, C.; Folgado, D.; Oliveira, J.; Franco, B.; Moreira, W.; Oliveira-Jr, A.; Cavaleiro, A.; Carvalho, R. TIP4.0: Industrial Internet of Things Platform for Predictive Maintenance. *Sensors* **2021**, *21*, 4676. [CrossRef] [PubMed]
4. Çınar, Z.M.; Abdussalam Nuhu, A.; Zeeshan, Q.; Korhan, O.; Asmael, M.; Safaei, B. Machine Learning in Predictive Maintenance towards Sustainable Smart Manufacturing in Industry 4.0. *Sustainability* **2020**, *12*, 8211. [CrossRef]
5. Taie, M.; Diab, M.; Elhelw, M. Remote Prognosis, Diagnosis and Maintenance for Automotive Architecture Based on Least Squares Support Vector Machine and Multiple Classifiers. In Proceedings of the 4th International Congress on Ultra Modern Telecommunications and Control Systems (ICUMT), St. Petersburg, Russia, 3–5 October 2012; pp. 128–134. [CrossRef]
6. Theissler, A. Detecting Known and Unknown Faults in Automotive Systems Using Ensemble-Based Anomaly Detection. *Knowl.-Based Syst.* **2017**, *123*, 163–173. [CrossRef]
7. Prytz, R.; Nowaczyk, S.; Rögnvaldsson, T.; Byttner, S. Predicting the Need for Vehicle Compressor Repairs Using Maintenance Records and Logged Vehicle Data. *Eng. Appl. Artif. Intell.* **2015**, *41*, 139–150. [CrossRef]
8. Ravi, A.; Surabhi, M.; Shah, C. Machine Learning Applications in Predictive Maintenance for Vehicles: Case Studies. *Int. J. Eng. Comput. Sci.* **2022**, *11*, 25628–25640. [CrossRef]
9. Singh, P.; Hungund, T.; Kukret, S. Revolutionizing System Operation and Maintenance in the Automobile Industry Through Machine Learning Applications. *Int. J. Comput. Sci. Inf. Technol.* **2024**, *16*, 93. [CrossRef]
10. Mafla-Yépez, C.N.; Morales-Bayetero, C.F.; Hernández-Rueda, E.P.; Benavides-Cevallos, I.B. Vehicle maintenance management based on machine learning in agricultural tractor engines. *DYNA* **2023**, *90*, 22–28. [CrossRef]
11. de Villiers, P.-R.H.; Jooste, J.L.; Lucke, D. Smart Maintenance System for Inner City Public Bus Services. *Procedia CIRP* **2023**, *120*, 285–290. [CrossRef]
12. Werbińska-Wojciechowska, S.; Giel, R.; Winiarska, K. Digital Twin Approach for Operation and Maintenance of Transportation System—Systematic Review. *Sensors* **2024**, *24*, 6069. [CrossRef] [PubMed]
13. Byttner, S.; Rögnvaldsson, T.; Svensson, M. Consensus Self-Organized Models for Fault Detection (COSMO). *Eng. Appl. Artif. Intell.* **2011**, *24*, 833–839. [CrossRef]
14. Killeen, P.; Ding, B.; Kiringa, I.; Yeap, T. IoT-Based Predictive Maintenance for Fleet Management. *Procedia Comput. Sci.* **2019**, *151*, 607–613. [CrossRef]
15. Silva, D. Detección de Anomalías en la Monitorización de una Flota de Autobuses. Master's Thesis, Universitat Politècnica de València, Valencia, Spain, 2022.
16. Massaro, A.; Selicato, S.; Galiano, A. Predictive Maintenance of Bus Fleet by Intelligent Smart Electronic Board Implementing Artificial Intelligence. *IoT* **2020**, *2*, 180–197. [CrossRef]
17. Bhave, D.; Adiga, D.T.; Powar, N.; Mckinley, T. Remaining Useful Life Prediction of Turbo Actuators for Predictive Maintenance of Diesel Engines. *PHM Soc. Eur. Conf.* **2021**, *6*, 11. [CrossRef]
18. Ridley, M. Explainable Artificial Intelligence (XAI). *Inf. Technol. Libr.* **2022**, *41*. [CrossRef]
19. Dwivedi, R.; Dave, D.; Naik, H.; Singhal, S.; Omer, R.; Patel, P.; Qian, B.; Wen, Z.; Shah, T.; Morgan, G.; et al. Explainable AI (XAI): Core Ideas, Techniques, and Solutions. *ACM Comput. Surv.* **2023**, *55*, 194. [CrossRef]
20. Gawde, S.; Patil, S.; Kumar, S.; Kamat, P.; Kotecha, K.; Alfarhood, S. Explainable Predictive Maintenance of Rotating Machines Using LIME, SHAP, PDP, ICE. *IEEE Access* **2024**, *12*, 29345–29361. [CrossRef]
21. Kobayashi, K.; Alam, S.B. Explainable, interpretable, and trustworthy AI for an intelligent digital twin: A case study on remaining useful life. *Eng. Appl. Artif. Intell.* **2024**, *129*, 107620. [CrossRef]
22. Moosavi, S.; Farajzadeh-Zanjani, M.; Razavi-Far, R.; Palade, V.; Saif, M. Explainable AI in Manufacturing and Industrial Cyber–Physical Systems: A Survey. *Electronics* **2024**, *13*, 3497. [CrossRef]
23. Mota, B.; Faria, P.; Corchado, J.; Ramos, C. Explainable Artificial Intelligence Applied to Predictive Maintenance: Comparison of Post-Hoc Explainability Techniques. In *Explainable Artificial Intelligence, xAI 2024, Communications in Computer and Information Science*; Longo, L., Lapuschkin, S., Seifert, C., Eds.; Springer: Cham, Switzerland, 2024; Volume 2156, pp. 318–335. [CrossRef]
24. McKinley, T.; Somwanshi, M.; Bhave, D.; Verma, S. Identifying NOx Sensor Failure for Predictive Maintenance of Diesel Engines Using Explainable AI. *PHM Soc. Eur. Conf.* **2020**, *5*, 11. [CrossRef]
25. Pashami, S.; Nowaczyk, S.; Fan, Y.; Jakubowski, J.; Paiva, N.; Davari, N.; Bobek, S.; Jamshidi, P.; Sarmadi, H.; Alabdallah, A.; et al. Explainable Predictive Maintenance. *arXiv* **2023**, arXiv:2306.05120. [CrossRef]
26. Amer, S.; Mohamed, H.K.; Mansour, M.B.M. Predictive Maintenance by Machine Learning Methods. In Proceedings of the 2023 Eleventh International Conference on Intelligent Computing and Information Systems (ICICIS), Cairo, Egypt, 21–23 November 2023; IEEE: Piscataway, NJ, USA, 2023; pp. 58–66. [CrossRef]

27. Hung, Y.-H. Developing an Improved Ensemble Learning Approach for Predictive Maintenance in the Textile Manufacturing Process. *Sensors* **2022**, *22*, 9065. [CrossRef]
28. Shivakarthik, S.; Karuppusamy, P.; Perikos, I.; Shi, F.; Nguyen, T.N. Maintenance of Automobiles by Predicting System Fault Severity Using Machine Learning. In *Sustainable Communication Networks and Application, Lecture Notes on Data Engineering and Communications Technologies*; Karuppusamy, P., Perikos, I., Shi, F., Nguyen, T.N., Eds.; Springer: Singapore, 2021; Volume 55, pp. 255–267. [CrossRef]
29. Wang, Y. A Comparative Analysis of Model Agnostic Techniques for Explainable Artificial Intelligence. *Res. Rep. Comput. Sci.* **2024**, *3*, 25–33. [CrossRef]
30. Kedar, M. Exploring the Effectiveness of SHAP over Other Explainable AI Methods. *Int. J. Sci. Res. Eng. Manag.* **2024**, *8*. [CrossRef]