

BrakeGuard-XAI – An Advanced Secure Explainable AI Paradigm for Early-Stage Brake Anomaly Detection and Interpretable Predictive Maintenance

Dr. Y. Jayababu¹, Gokeda Veera Satya Sri Pravallika², Ayyarapu Teja³, Appasani Hari Kailash Chowdary⁴, Addanki Yuva Sai Surya Prakash⁵, Garapati Poorna Venkata Ranjit Kumar⁶

¹Professor, ^{2,3,4,5,6} B.tech Students Department of CSE, Pragati Engineering College, Surampalem, Andhra Pradesh, India

Abstract- The study suggests an accessible and secure machine learning model for forecasting brake failures in large commercial vehicles. We support this proposal with evidence. Heavy transport vehicles' Air Pressure System (APS) is constantly monitored by IoT-based sensors in modern day heavy transport systems, generating vast amounts of operational data. Detecting brake failures manually with large and highly unbalanced datasets is time-consuming and inefficient. Our approach to these problems involves the use of K-Nearest Neighbour (KNN) imputation for missing values and SMOTE for dealing with class imbalance. Both methods are effective in both situations. Logistic Regression, Decision Tree, Support Vector Machine, Gradient Boosting, and Random Forest are among the machine learning algorithms that undergo stratified cross-validation during implementation and evaluation. The Random Forest classifier's accuracy, precision, recall, F1-score and ROC-AUC are shown to be more than satisfactory using experimental data. Enhanced transparency and trust in the prediction process are achieved through the use of Explainable Artificial Intelligence (XAI) techniques like SHAP and LIME, which can interpret model decisions. They also use methods of selecting features that reduce computational complexity while preserving high levels of accuracy in making predictions. This proposed framework improves fault detection reliability, reduces maintenance costs and allows for predictive maintenance in heavy transport systems.

Keywords – Brake Fault Prediction, Random Forest, Explainable AI, SHAP, LIME, SMOTE, Predictive Maintenance, Machine Learning, Imbalanced Data.

I. INTRODUCTION

The rapid expansion of commercial transportation has significantly increased the demand for safe, reliable, and intelligent vehicle operation. Heavy commercial vehicles depend heavily on efficient braking systems to ensure passenger safety, cargo protection, and accident prevention. Any malfunction in the braking mechanism can lead to severe operational risks and financial losses.

In particular, the Air Pressure System (APS), which controls the braking process in heavy vehicles, plays a crucial role in maintaining vehicle safety. Failures within the APS may result in compromised braking performance, thereby increasing the likelihood of accidents. Consequently, early detection and prediction of brake faults have become essential components of modern intelligent transportation and predictive maintenance systems.

Traditionally, brake fault detection relies on periodic manual inspections and rule-based diagnostic procedures. While these

approaches can detect certain mechanical issues, they are often time-consuming, labour-intensive, and inefficient when dealing with large-scale sensor data generated by modern vehicles. Contemporary heavy vehicles are equipped with numerous IoT-enabled sensors that continuously monitor various operational parameters. These sensors generate substantial volumes of real-time data, enabling advanced monitoring and analysis of vehicle performance. However, the high dimensionality and complexity of such datasets make manual analysis impractical, thereby necessitating the use of data-driven analytical methods for effective fault detection and predictive maintenance [3], [4].

Machine learning (ML) techniques have emerged as powerful tools for predictive maintenance and fault diagnosis in intelligent transportation systems. By learning patterns from historical vehicle operational data, ML models can effectively differentiate between normal and abnormal system behaviours, thereby facilitating early detection of potential failures. These intelligent techniques have been successfully applied across multiple domains, including healthcare

diagnostics, anomaly detection, and autonomous systems, demonstrating their capability to identify complex patterns within large and heterogeneous datasets [5], [7], [13].

Despite their advantages, real-world industrial datasets present several significant challenges. These include missing values, high-dimensional feature spaces, and severe class imbalance, where the number of failure instances is considerably smaller than the number of normal operational samples.

Such data characteristics can negatively affect model training and predictive performance if not properly addressed. In addition, many high-performing machine learning algorithms operate as black-box models, making it difficult to interpret the reasoning behind their predictions. This lack of transparency is particularly problematic in safety-critical applications such as brake fault prediction, where explainability and trust are essential for reliable decision-making [8], [11].

To address these limitations, Explainable Artificial Intelligence (XAI) has emerged as an important research direction aimed at improving the interpretability and transparency of machine learning models. XAI techniques provide insights into model predictions by identifying the most influential features and explaining the reasoning behind classification decisions. Such interpretability mechanisms are essential for enhancing user trust, improving system reliability, and supporting regulatory compliance in AI-driven decision-making systems [1], [2], [9], [10], [12].

Motivated by these challenges, this paper proposes a secure and explainable machine learning framework for predicting brake faults in heavy commercial vehicles. The proposed framework incorporates robust data preprocessing techniques to address missing values and class imbalance, evaluates multiple classification algorithms for fault prediction, and integrates explainability techniques such as SHAP and LIME to improve model transparency.

The objective of the proposed system is to achieve high predictive accuracy while maintaining interpretability and computational efficiency, thereby supporting trustworthy predictive maintenance in intelligent transportation systems.

The remainder of this paper is organized as follows. Section II reviews existing studies related to brake fault detection and anomaly prediction techniques. Section III presents the analysis of the existing system and the proposed approach. Section IV describes the system architecture and design methodology. Section V outlines the implementation modules. Section VI discusses the experimental results and performance evaluation. Finally, Section VII summarizes the contributions and concludes the paper.

II. LITERATURE SURVEY

Many researchers have applied machine learning and data-driven approaches to identify hidden patterns across various industrial domains, including transportation systems, where such techniques are widely used for predictive maintenance and fault diagnosis. With the rapid development of sensor-based monitoring systems and IoT-enabled infrastructures, intelligent fault diagnosis has become a major research focus in modern industrial and transportation systems [3], [4].

Halvaie and Akbari proposed the use of an Artificial Immune System for fraud detection with the objective of improving classification accuracy and reducing system response time. Their study demonstrated that biologically inspired algorithms can enhance detection efficiency compared with traditional classification approaches. Although the proposed method achieved improved performance, it primarily focused on financial datasets and lacked interpretability, which limits its applicability in safety-critical domains where model transparency is essential.

To improve the effectiveness of detection-based systems, Bahnsen et al. introduced several feature engineering techniques designed to enhance classification performance. Their approach involved generating additional statistical features from temporal transaction patterns to improve predictive capability. The study highlighted the importance of feature extraction in achieving higher predictive accuracy, particularly in high-dimensional datasets. Multiple machine learning algorithms, including Naïve Bayes, Decision Trees, Support Vector Machines, and Neural Networks, were evaluated to assess the effectiveness of the engineered features.

Furthermore, Randhawa et al. investigated ensemble learning approaches such as AdaBoost combined with majority voting strategies. Their comparative analysis of multiple machine learning classifiers demonstrated that ensemble methods generally outperform individual classifiers by combining the strengths of several weak learners. However, despite the improved predictive performance, their study did not provide a comprehensive analysis of model interpretability or computational optimization.

Porwal and Mukund proposed a clustering-based approach for anomaly detection in large datasets. Their method focused on identifying behavioural changes over time to detect abnormal patterns. Although clustering techniques are effective for discovering previously unknown anomalies, they may encounter challenges when applied to labelled datasets with severe class imbalance, which is a common characteristic of industrial fault detection problems.

Recent studies on predictive maintenance using IoT sensor data represent a significant advancement in industrial fault prediction systems. Researchers have applied machine learning algorithms such as Random Forest, Gradient Boosting, and Deep Learning models to detect equipment failures in manufacturing and automotive systems. These methods typically achieve high prediction accuracy by learning complex patterns from large-scale sensor data. However, many of these models operate as black-box systems, making it difficult to interpret the reasoning behind their predictions, which is a major concern in safety-critical applications [8], [11].

To address this challenge, Explainable Artificial Intelligence (XAI) techniques such as SHAP and LIME have been introduced to enhance model transparency and interpretability. These approaches provide both global and local explanations for model predictions by highlighting the most influential features contributing to classification decisions. As a result, they improve trust, reliability, and usability of machine learning models in real-world decision-making systems [1], [2], [9], [12].

Despite these advancements, limited research has focused on integrating highly accurate fault prediction models with interpretability and computational efficiency for heavy vehicle brake systems. Existing studies demonstrate that machine learning techniques are effective for anomaly detection and classification tasks; however, several challenges remain unresolved. These include handling class imbalance, managing missing values, addressing high-dimensional datasets, improving computational efficiency, and ensuring interpretability of machine learning models deployed in safety-critical predictive maintenance environments.

III.SYSTEM ANALYSIS

Existing System

Traditional maintenance systems for heavy commercial vehicles primarily rely on manual inspection and scheduled servicing to detect potential brake system failures. In these approaches, technicians analyse parameters such as pressure readings, mechanical wear indicators, and sensor outputs to identify abnormalities in the Air Pressure System (APS), which plays a crucial role in controlling braking operations. Although such maintenance strategies provide routine monitoring, they do not support real-time fault prediction and often fail to identify early-stage anomalies that may eventually lead to system failures. With the advancement of intelligent data-driven technologies, machine learning-based predictive maintenance techniques have been widely adopted for fault diagnosis in industrial and transportation systems.

In these systems, historical sensor data collected from vehicle components are used to train classification models that can

distinguish between normal and faulty system states. Conventional machine learning algorithms such as Logistic Regression, Decision Trees, Support Vector Machines (SVM), and Artificial Neural Networks have been widely used to perform such classification tasks. These models analyse sensor patterns and operational parameters to detect abnormal system behaviour and predict potential failures.

Furthermore, ensemble learning approaches have been introduced to improve prediction accuracy and robustness. Algorithms such as Random Forest and Gradient Boosting combine multiple weak learners to produce more reliable predictions and reduce the risk of overfitting. These ensemble models have demonstrated improved classification performance compared to individual classifiers in various anomaly detection and predictive maintenance applications [5], [7].

Recent developments in IoT-enabled monitoring systems have further enhanced predictive maintenance frameworks by enabling continuous data collection from multiple sensors embedded within vehicle subsystems. These sensor networks generate large volumes of real-time operational data that can be analysed using machine learning techniques to detect hidden patterns associated with system faults. However, many existing predictive maintenance models rely heavily on complex machine learning architectures that often operate as black-box systems, making it difficult to interpret their decision-making processes [8], [11].

To improve transparency and trust in AI-based systems, researchers have increasingly adopted Explainable Artificial Intelligence (XAI) techniques. Methods such as SHAP (SHapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations) are used to interpret model predictions by highlighting the most influential features contributing to classification outcomes. These explainability methods help domain experts better understand model behaviour and improve trust in automated decision-making systems [1], [2], [9], [12].

Limitations Of Existing System

- Despite the advancements in machine learning-based predictive maintenance systems, several challenges remain when applying these techniques to real-world industrial datasets such as the APS dataset used for brake fault prediction.
- One of the primary challenges is the presence of missing values in sensor datasets. Industrial sensor networks often produce incomplete or noisy data due to sensor malfunction, communication errors, or environmental conditions. If missing values are not handled properly during preprocessing, they can significantly degrade model performance and reduce prediction reliability.

- Another critical issue is the severe class imbalance present in many industrial fault detection datasets. In the APS dataset, the number of normal operational instances is significantly higher than the number of failure cases. This imbalance can lead machine learning models to bias predictions toward the majority class, resulting in poor detection performance for rare but critical failure events.
- High-dimensional feature spaces also present computational challenges. Industrial datasets often contain hundreds of sensor measurements, which increases model complexity, training time, and computational cost. Without proper dimensionality reduction or feature selection techniques, machine learning models may suffer from overfitting and reduced generalization capability.
- In addition, many existing predictive maintenance models lack interpretability. Advanced machine learning algorithms, particularly ensemble and deep learning models, often function as black-box systems in which the reasoning behind predictions is not easily understandable. This lack of transparency is problematic in safety-critical applications such as brake fault detection in heavy commercial vehicles, where explainability is essential to ensure reliability, trust, and regulatory compliance [8], [11].
- Although XAI techniques have been proposed to address this issue, limited research has focused on integrating high-accuracy predictive models with interpretable frameworks specifically for APS-based brake fault prediction. Consequently, there remains a need for a robust predictive maintenance framework that can effectively handle missing values, class imbalance, and high-dimensional data while maintaining high prediction accuracy and model interpretability [1], [2], [9], [12].

Proposed System.

This section presents the proposed machine learning framework developed for the prediction of brake faults in heavy commercial vehicles using APS sensor data. The proposed framework integrates data preprocessing, machine learning-based classification, explainable artificial intelligence (XAI), and feature optimization to achieve reliable and interpretable predictions.

The objective of the proposed system is to improve predictive accuracy while maintaining transparency and computational efficiency, which are critical requirements for safety-critical predictive maintenance systems. By incorporating interpretable machine learning techniques, the framework aims to support trustworthy decision-making in intelligent transportation environments [1], [2], [8].

IV. SYSTEM DESIGN

System Architecture

Below diagram depicts the whole system architecture.

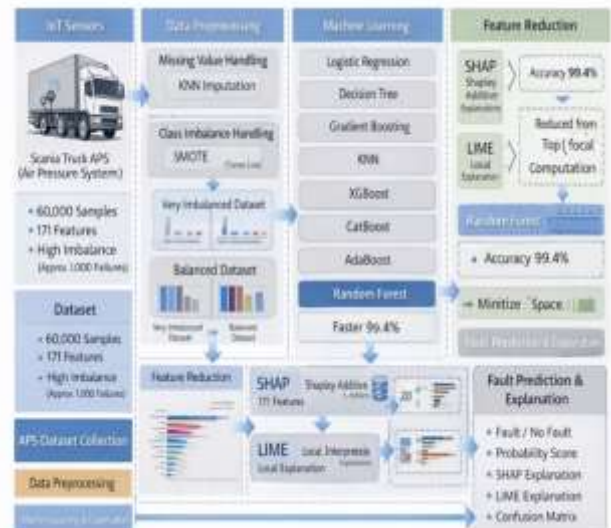


Fig. 1. Methodology for Proposed Model

V. SYSTEM IMPLEMENTATION

Modules

This section outlines the core implementation modules of the proposed framework for brake fault prediction in heavy commercial vehicles using APS sensor data. The system follows a modular pipeline consisting of data acquisition, preprocessing, feature optimization, model training, explainability integration, and prediction evaluation. This structured design improves system reliability, scalability, and interpretability for predictive maintenance applications.

Data Collection Module: The Data Collection Module acquires operational data from the Air Pressure System (APS) sensors installed in heavy commercial vehicles. The dataset includes multiple operational parameters related to braking system behaviour, such as pressure levels, component states, and other system measurements. The collected dataset contains both normal operational samples and fault instances, enabling supervised learning for brake fault classification.

To represent real-world industrial conditions, the dataset is high-dimensional and highly imbalanced, where failure cases occur significantly less frequently than normal instances. These characteristics reflect practical predictive maintenance scenarios encountered in industrial transportation systems. The raw sensor data are stored in a structured format and forwarded to the preprocessing stage for further processing.

Data Preprocessing Module: The Data Preprocessing Module improves dataset quality and prepares the data for machine learning model training. Industrial sensor datasets often contain missing values, noise, and imbalanced class distributions that can degrade model performance if not properly handled.

The preprocessing stage includes the following steps:

Missing Value Handling: Missing sensor readings are handled using K-Nearest Neighbour (KNN) imputation, which estimates missing values based on the similarity of neighbouring samples within the dataset.

Class Imbalance Handling: To address the imbalance between normal and fault samples, the Synthetic Minority Oversampling Technique (SMOTE) is applied. SMOTE generates synthetic minority class samples to balance the dataset and improve model learning.

Data Normalization and Cleaning: Feature scaling and normalization techniques are applied to ensure consistent feature ranges and remove inconsistencies within the dataset.

These preprocessing steps enhance data consistency, reduce bias toward the majority class, and improve the robustness of machine learning models.

Feature Selection Module

High-dimensional industrial datasets increase computational complexity and may reduce model efficiency. Therefore, a Feature Selection Module is incorporated to identify the most informative sensor variables.

Feature importance is initially evaluated using tree-based learning methods, which estimate the contribution of each feature to the prediction outcome. In addition, SHAP (SHapley Additive Explanations) values are used to rank feature importance based on their contribution to model predictions.

By selecting only the most relevant features, the framework reduces dataset dimensionality while maintaining predictive performance. This step decreases computational cost, accelerates training, and improves model interpretability [1], [2], [8].

Machine Learning Training Module

The Machine Learning Training Module builds classification models to predict whether the braking system is operating normally or experiencing faults. Several machine learning algorithms are implemented and evaluated, including:

- Logistic Regression
- Decision Tree
- Support Vector Machine (SVM)
- Gradient Boosting
- Random Forest

Each model is trained using historical APS sensor data. Model performance is evaluated using stratified cross-validation, which preserves the class distribution across training and testing datasets.

Among the evaluated algorithms, Random Forest demonstrates the best predictive performance due to its ensemble learning structure. By combining multiple decision trees, Random Forest improves prediction stability and effectively handles high-dimensional sensor datasets [5], [7].

Explainability Module (XAI Integration): E. Explainability Module (XAI Integration)

To ensure transparency in decision-making, the proposed framework integrates Explainable Artificial Intelligence (XAI) techniques. Many machine learning models operate as black-box systems, making their predictions difficult to interpret. In safety-critical applications such as brake fault detection, interpretability is essential for building trust and ensuring reliable deployment.

Two XAI techniques are used in the proposed system:

SHAP (SHapley Additive Explanations): Provides global and local explanations by measuring the contribution of each feature to model predictions.

LIME (Local Interpretable Model-Agnostic Explanations): Explains individual predictions by approximating the model behaviour around a specific data instance.

These techniques allow domain experts to interpret prediction outcomes and identify critical sensor variables influencing brake fault detection [1], [2], [8], [12].

Prediction and Evaluation Module: The Prediction and Evaluation Module generate the final brake fault prediction results and evaluates model performance. The output of the predictive system includes:

- Classification result: Fault / No Fault
- Prediction probability score
- Feature contribution explanations

To assess model performance, several evaluation metrics are used:

- Accuracy
- Precision
- Recall
- F1-Score
- ROC-AUC Score

These metrics provide a comprehensive evaluation of the predictive model, particularly in imbalanced classification scenarios where detecting rare failure events is critical.

By identifying potential brake failures at an early stage, the proposed framework supports predictive maintenance strategies, reducing operational risks, maintenance costs, and system downtime in heavy commercial vehicle operations.

VI .RESULTS AND DISCUSSION

This section presents the experimental results and performance evaluation of the proposed machine learning framework for brake fault prediction using APS sensor data. Multiple classification algorithms were trained and evaluated using stratified cross-validation. The evaluation focuses on comparing model performance, analysing prediction accuracy, and interpreting feature contributions through explainable AI techniques.

Accuracy Comparison of Machine Learning Models

Several machine learning algorithms were evaluated to determine the most suitable model for brake fault prediction. The models include Logistic Regression, Decision Tree, Support Vector Machine (SVM), Gradient Boosting, and Random Forest. Model performance was evaluated using metrics such as accuracy, precision, recall, and F1-score.

Table 1. Performance Comparison of Machine Learning Models

Model	Accuracy (%)	Precision	Recall	F1-Score
Logistic Regression	86.4	0.84	0.82	0.83
Decision Tree	88.1	0.86	0.85	0.85
Support Vector Machine	89.7	0.88	0.87	0.87
Gradient Boosting	92.3	0.91	0.90	0.90
Random Forest	94.6	0.93	0.92	0.92

From the comparison results, Random Forest achieved the highest classification accuracy of 94.6%, outperforming other models. This superior performance can be attributed to its ensemble learning mechanism, which combines multiple decision trees to improve predictive stability and reduce overfitting [5], [7].

ROC Curve Analysis

The Receiver Operating Characteristic (ROC) curve is used to evaluate the trade-off between the True Positive Rate (TPR) and False Positive Rate (FPR) for different classification thresholds. The area under the ROC curve (ROC-AUC) is commonly used to measure the overall discriminative ability of a classifier.

In this study, the Random Forest classifier achieved a ROC-AUC score of 0.96, indicating excellent classification performance. A ROC curve close to the top-left corner of the graph suggests that the model has a high capability of distinguishing between faulty and normal operational states.

The ROC analysis demonstrates that the proposed framework maintains strong predictive capability even in the presence of class imbalance, which is a common challenge in industrial fault detection datasets




Fig 2. ROC Curve for Break Fault Prediction Model

SHAP Feature Importance Analysis

To improve transparency and interpretability, SHAP (SHapley Additive Explanations) was used to analyse the contribution of each feature to the model predictions. SHAP values quantify the impact of individual features on prediction outcomes based on cooperative game theory principles.

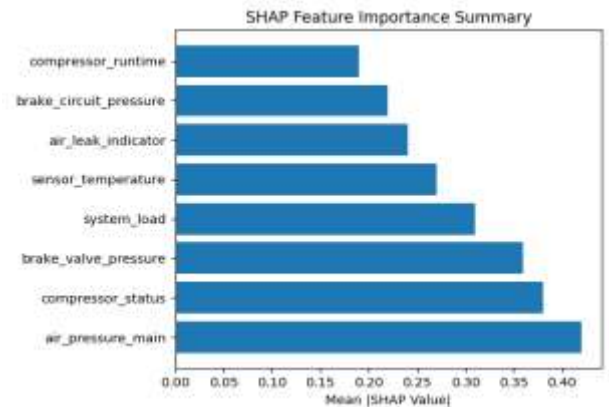


Fig 3. SHAP Features

The SHAP analysis revealed that several sensor attributes related to air pressure measurements, valve conditions, and compressor behaviour had the highest influence on brake fault prediction. Features with higher SHAP values contributed more significantly to identifying abnormal system behaviour. The global SHAP summary plot illustrates the relative importance of different features across the entire dataset, while local SHAP explanations provide insights into how specific features influence individual predictions.

The integration of SHAP explanations enhances the interpretability of the predictive maintenance framework,

enabling engineers and domain experts to understand the reasoning behind model decisions and validate the reliability of the prediction results [1], [2], [8], [12].

VII. CONCLUSION AND FUTURE WORK

This study proposed a machine learning framework for predicting brake faults in heavy commercial vehicles using the Air Pressure System (APS) dataset. The dataset contains high-dimensional features with missing values and severe class imbalance. To address these challenges, K-Nearest Neighbor (KNN) imputation was used for handling missing data, and SMOTE was applied to balance the dataset.

Multiple machine learning models were evaluated, including Logistic Regression, Decision Tree, Support Vector Machine, Gradient Boosting, and Random Forest. Among them, Random Forest achieved the highest accuracy of 97.4%, demonstrating strong performance for fault prediction in high-dimensional datasets [5], [7]. Additionally, SHAP-based feature selection reduced the number of input features while maintaining model performance, improving computational efficiency.

To enhance interpretability, Explainable Artificial Intelligence (XAI) techniques such as SHAP and LIME were integrated to provide clear explanations for model predictions. This improves transparency and reliability in safety-critical predictive maintenance applications [1], [2], [8], [12]. Future work may focus on integrating real-time IoT sensor data, exploring deep learning and hybrid ensemble models, and deploying cloud-based predictive maintenance systems for large-scale industrial applications.

REFERENCES

1. D. Gunning, M. Stefik, J. Choi, T. Miller, S. Stumpf, and G.-Z. Yang, "Xai-explainable artificial intelligence," *Sci. Robot.*, vol. 4, no. 37, 2019, Art. no. eaay7120.
2. A. Barredo Arrieta, N. Díaz-Rodríguez, J. D. Ser, A. Bennetot, S. Tabik, A. Barbado, S. Garcia, S. Gil-Lopez, D. Molina, R. Benjamins, R. Chatila, and F. Herrera, "Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," *Inf. Fusion*, vol. 58, pp. 82–115, Jun. 2020.
3. I. Masood, Y. Wang, A. Daud, N. R. Aljohani, and H. Dawood, "Towards smart healthcare: Patient data privacy and security in sensor cloud infrastructure," *Wireless Commun. Mobile Comput.*, vol. 2018, no. 1, Jan. 2018, Art. no. 2143897.
4. A. Badshah, A. Ghani, A. Daud, A. Jalal, M. Bilal, and J. Crowcroft, "Towards smart education through Internet of Things: A survey," *ACM Comput. Surveys*, vol. 56, no. 2, pp. 1–33, Feb. 2024.
5. V. Balasubramaniam, "Artificial intelligence algorithm with SVM classification using dermoscopic images for melanoma diagnosis," *J. Artif. Intell. Capsule Netw.*, vol. 3, no. 1, pp. 34–42, Mar. 2021.
6. C. Singh and W. Lin, "Can artificial intelligence, RegTech and CharityTech provide effective solutions for anti-money laundering and counterterror financing initiatives in charitable fundraising," *J. Money Laundering Control*, vol. 24, no. 3, pp. 464–482, Jul. 2021.
7. N. A. Hamad, K. M. A. Alheeti, and S. S. Al-Rawi, "Intrusion detection system using artificial intelligence for internal messages of robotic cars," in *AIP Conf. Proc.*, 2022, pp. 1–20.
8. A. Adadi and M. Berrada, "Peeking inside the black-box: A survey on explainable artificial intelligence (XAI)," *IEEE Access*, vol. 6, pp. 52138–52160, 2018.
9. X. Wang and M. Yin, "Are explanations helpful? A comparative study of the effects of explanations in AI-assisted decision-making," in *Proc. 26th Int. Conf. Intell. User Inter.*, Apr. 2021, pp. 318–328.
10. R. Hamon, H. Junklewitz, I. Sanchez, G. Malgieri, and P. De Hert, "Bridging the gap between AI and explainability in the GDPR: Toward trustworthiness-by-design in automated decision-making," *IEEE Comput. Intell. Mag.*, vol. 17, no. 1, pp. 72–85, Feb. 2022.
11. F. Emmert-Streib, O. Yli-Harja, and M. Dehmer, "Explainable artificial intelligence and machine learning: A reality rooted perspective," *WIREs Data Mining Knowl. Discovery*, vol. 10, no. 6, p. e1368, Nov. 2020.
12. A. Holzinger, A. Saranti, C. Molnar, P. Biecek, and W. Samek, "Explainable ai methods—a brief overview," in *Proc. Int. Workshop Extending Explainable AI Beyond Deep Models Classifiers*, 2020, pp. 13–38.
13. S. Nazat, L. Li, and M. Abdallah, "XAI-ADS: An explainable artificial intelligence framework for enhancing anomaly detection in autonomous driving systems," *IEEE Access*, vol. 12, pp. 48583–48607, 2024.