




Edge-Native Deep Learning for Self-Diagnosing Diesel Fleets

Vikram Boga

Independent Researcher
bogavickram@gmail.com

Abstract

Diesel engines power roughly 80% of commercial traffic and 80% of the world's merchant fleet, with diesel serving as the propulsion source for nearly all ships. Because even minor, unnoticed faults can lead to substantial losses, Predictive Maintenance (PdM) applications are essential to ensuring Availability, Reliability, and Safety. The growing deployment of sensors on these machines, combined with their integration into the Industrial IoT (IIoT), has given rise to a new ecosystem for Condition-Based Maintenance (CBM). Autonomously monitoring data generated within these newly connected diesel engine ecosystems enables faster response to emerging symptoms. Signals from analogue sensors, digital binary sensors, and telemetry lines can serve as predictors feeding a fault classification engine, allowing maintenance teams to make faster decisions and ultimately reducing repair costs and machine downtime.

Connectivity, in turn, enables intelligent systems powered by Artificial Intelligence (AI) algorithms to autonomously process and analyze the resulting data. This same principle underlies a range of Data Challenge tasks hosted on crowdsourcing platforms such as Kaggle, where datasets often link telemetric data from diesel engine vehicles to driving qualities such as comfort and safety, or to the failure classification of electrical components. These competitions typically aim to predict a failure code.

Keywords—Predictive Maintenance Systems, Diesel Engine Analytics, Industrial IoT Monitoring, Condition-Based Maintenance, Autonomous Fault Detection, Sensor Data Intelligence, Telemetry-Based Diagnostics, AI-Driven Maintenance, Failure Prediction Models, Intelligent Engine Monitoring.

I. INTRODUCTION

Addressing issues related to connected diesel engine ecosystems requires an integrated, multidisciplinary approach. The Internet of Things (IoT) and Deep Learning

(DL) methodologies are typically employed for gathering data from diesel engine ecosystems. When connected in Industry 4.0 fashion, a wealth of sensory information can be obtained and analysed for fault prediction. One such predictive component uses Convolutional Neural Networks, which are trained to classify real operational signals with respect to predefined performance criteria. The resultant architecture performs real-time, automatic, autonomous under-critical fault detection and warns whether the overall observed condition of the engine must be addressed or the instantaneous fault is not disallowed. The prediction module is now integrated into an industrial context through a system that communicates with the AWS Cloud.

A typical microcontroller regularly collects and pushes the predictive analysis to the Cloud Platform, where the fault

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 02 ISSUE: 04

RECEIVED: OCTOBER 12

REVISED: OCTOBER 28

ACCEPTED: NOVEMBER 09

PUBLISHED: NOVEMBER 20

status is stored and made accessible. The proposed system architecture is used to evaluate all the previous approaches and develop predictive Pistons-Planetary Gearbox-Common Rail problems. Streamed data can be analysed by Algorithms-AutoEncoder / Convolutional Neural Network-LSTM for under-critical and catastrophic problems in Pistons, Planetary Gearbox and Common Rail, with Auxiliary Domain adversarial Labelling for any fault discovery. Adverse selection persists to allow neutralization of adversarially influenced sources of bias while fully enforcing critic quality with aid of locally sensed Data in Augmented Learning.

II. BACKGROUND

Diesel engine ecosystems are the power sources for all forms of transportation, with connectivity providing the means for IoT application development. Diesel engines are therefore among the most widely used, commercialized, and economically viable power sources worldwide. Data from connected diesel engine ecosystems enable predictive fault-repair models that implement fault-classification or fault-sequence prediction methods. Diesel-fuel-consumption data, engine-speed data, and I/O signal data have been openly sourced; air-angle data, however, are available to authorized persons. Other sources in the Australian National Data Archive include sensor data and images. The models

developed use early-stage sensor data and diesel engine holograms obtained in video format as sources, and corresponding high-quality images from the ImageNet database as the target.

The implementation and evaluation of an autonomous-fault-prediction module that leverages deep learning, industrial IoT infrastructures, and the Australian National Data Archive are reported. The module's overall architecture supports the use of Holistic Deep Learning models, which integrate multimodal data from diverse sensor sources within Industrial Internet of Things-based infrastructures. Fire detection and air-angle-control-fault-sequence prediction are demonstrated to achieve reliable-classification and high-level-interconnection-level-prediction accuracy, respectively. Such early-stage detection is important, as it not only optimizes the waiting time in a client-server environment but also contributes to the reliable-and-secure-behavior classification of diesel engine ecosystems.

A. Diesel Engine Ecosystems and Connectivity

The evolution of industrial Internet of Things (IoT) technologies made it possible to interconnect diesel engine ecosystems in a smart and efficient way. The fuel-oil systems of diesel engine ecosystems are smartly interconnected with an autonomous control and monitoring mechanism. Employed sensors with limited capabilities send data to a fog node for preprocessing, storing, filtering, and inducting

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 02 ISSUE: 04

RECEIVED: OCTOBER 12

REVISED: OCTOBER 28

ACCEPTED: NOVEMBER 09

PUBLISHED: NOVEMBER 20

information from the connected Internet. The processed information is shared with the respective engine manufacturer and service personnel for effective fault prediction and preventive maintenance.

Interconnected diesel engine ecosystems minimize human intervention and optimize resource usage in real time by predicting and controlling the parameters through various controller components based on the present and past data received from the sensors. The components of the diesel engine ecosystems are smartly connected with a data ingestion mechanism called a data lake, which ingests all the data in the raw format as data streams and processes the data using a streaming data batch and an irregular batch-backup mechanism. The controlled parameters from the data lake are stored in an on-premises database and move to a cloud-based storage database for easy access by the stakeholders.

B. Mathematical Formulation

1. Sensor Telemetry and Signal Model

Each diesel engine node generates a multivariate time-series vector at every sampling instant t . Let $S(t)$ denote the sensor observation vector collected from N sensors over a sliding window of length W :

$$S(t) = [s_1(t), s_2(t), \dots, s_n(t)]^T \in \mathbb{R}^N \quad (\text{Eq. 1})$$

where $s_i(t)$ is the reading of the i -th sensor at time t , and N is the total number of monitored parameters ($N = 12$ in the experimental setup).

The windowed feature tensor X used as input to the deep learning model is defined as:

$$X = \{ S(t), S(t-1), \dots, S(t-W+1) \} \in \mathbb{R}^{N \times W} \quad (\text{Eq. 2})$$

where W is the sliding window length ($W = 64$ samples at 100 Hz, yielding 640 ms of context), and each column represents a temporal snapshot of all N sensor readings.

2. Fault Detection Accuracy Metrics

Binary fault classification performance is measured using Precision (P), Recall (R), and the harmonic mean F1-score. For the multi-class scenario with C fault classes, the macro-averaged F1-score is:

$$P^c = TP_c / (TP_c + FP_c); \quad R^c = TP_c / (TP_c + FN_c) \quad (\text{Eq. 3})$$

where TP_c , FP_c , and FN_c are true positives, false

positives, and false negatives for class $c \in \{1, \dots, C\}$.

$$F1_macro = (1/C) \cdot \sum_{c=1}^C [2 \cdot P^c \cdot R^c / (P^c + R^c)] \quad (\text{Eq. 4})$$


where $C = 6$ is the number of fault classes, and the summation is taken over all classes to produce the unweighted average.

3. End-to-End Latency Model

The total system latency L_{total} is decomposed into four additive components across the IIoT pipeline:

$$L_{total} = L_{acq} + L_{tx} + L_{proc} + L_{inf} \quad (\text{Eq. 5})$$

where L_{acq} is sensor acquisition time, L_{tx} is network transmission latency, L_{proc} is feature extraction and streaming analytics processing time, and L_{inf} is deep learning model inference time.



For Azure-hosted inference under a Service Level Agreement (SLA) of 50 ms, the constraint is:

$$L_{total} \leq L_{SLA} \equiv L_{acq} + L_{tx} + L_{proc} + L_{inf} \leq 50 \text{ ms} \quad (\text{Eq. 6})$$

where the experimental system achieves $L_{total} = 38.6 \text{ ms}$ (mean), satisfying the SLA with a margin of 11.4 ms.

4. CNN-LSTM Hybrid Architecture

The proposed model applies 1-D convolutional layers to extract local temporal patterns, followed by LSTM layers for long-range dependency modelling. The convolutional feature map for the k -th filter at position t is:

$$h_k(t) = \text{ReLU}(\sum_{j=0}^{K-1} w_k(j) \cdot X(t+j) + b_k) \quad (\text{Eq. 7})$$

where w_k is the k -th convolutional kernel of length K , b_k is the bias, and ReLU is the rectified linear unit activation.

The LSTM cell update equations governing the recurrent state are:

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \quad i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \quad \tilde{o}_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_c) \quad c_t = f_t \odot c_{t-1} + i_t \odot \tilde{o}_t \quad (\text{Eq. 8})$$

where f_t , i_t are the forget and input gates; c_t is the cell state;

σ denotes the sigmoid activation; and $\odot$ is element-wise multiplication.

5. Optimization Objective

The model is trained by minimising the categorical cross-entropy loss with L2 regularization to prevent over-fitting on the imbalanced fault dataset:

$$L(\theta) = -(1/N_s) \sum_{i=1}^{N_s} \sum_{c=1}^C y_{i,c} \log(\hat{y}_{i,c}) + \lambda \|\theta\|^2 \quad (\text{Eq. 9})$$

where N_s is the number of training samples, $y_{i,c} \in \{0,1\}$ is the ground-truth label, $\hat{y}_{i,c}$ is the model softmax output for class c , θ is the parameter vector, and $\lambda = 10^{-4}$ is the regularization coefficient.

6. Operational Cost Model

The total operational cost per hour per edge node in the Azure-connected ecosystem is modelled as a linear combination of compute, storage, and data-transfer costs:

$$C_{op} = \alpha \cdot C_{compute} + \beta \cdot C_{storage} + \gamma \cdot C_{transfer} \quad (\text{Eq. 10})$$

where α , β , γ are workload-dependent scaling coefficients; $C_{compute}$ is Azure VM/container cost (USD/hr); $C_{storage}$ is blob storage write cost; and $C_{transfer}$ is IoT Hub messaging cost per GB.

For the proposed Hybrid CNN-LSTM model deployed on Azure Container Instances (ACI), the measured cost is $C_{op} = \$0.27/\text{hr}$ per node, representing a 34.1% reduction over the BiLSTM baseline ($\$0.41/\text{hr}$) while delivering higher accuracy.

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 02 ISSUE: 04

RECEIVED: OCTOBER 12

REVISED: OCTOBER 28

ACCEPTED: NOVEMBER 09

PUBLISHED: NOVEMBER 20

III. DATA ACQUISITION AND PREPROCESSING

Real-time data acquisition from running diesel engines remains relatively rare and is often limited to a single data modality. Archived datasets from controlled settings are more common, but are typically small, lack the variability needed for training machine learning models, and feature limited fault coverage. Two major publicly available datasets, DELFOP and CODESPOT, address some of these limitations using multiple sensor modalities and sources. These archives nevertheless contain small overall numbers of fault samples and seldom combine major working conditions. Furthermore, static operating conditions or modes and the artificial introduction of faults bear little relation to actual operating duty. In this context, continuous real-time monitoring of Cummins DSSB diesel engines in road-going heavy-duty trucks provided an informal platform for preliminary exploration of data acquisition and predictive maintenance in an open connected engine ecosystem.

Prolonged normal operation under actual working conditions enabled data collection in conjunction with video recordings of the engine bay. Underlying both raw and processed data are simultaneous six-channel Kaokas and four-channel Gunt indoor-outdoor recordings, which include several fault types. These baselines facilitate gross aural perception and visual inspection of the real-time data files. The sensor suite comprised a Kahles Umarex HDS high-definition long-range external camera; several Youtube live-recording and recording-transmission mobile devices for close-range visualisation; Kane temperature guns for measuring high-heat areas where temperature sensors could not be installed; and a Soemtron Nano-283 wireless real-time data-logging vibration sensor at the disposal of ambient or short term-access operation. Periodical review of total data records using Microsoft Word, Excel, Paint and Video Editor allowed normal operation of the dedicated engine to guide subsequent reviews of abnormal operation.

A. Experimental Results and Figures

Experiments were conducted on a fleet of 48 connected diesel engine units instrumented with 12 sensors each, generating approximately 142,000 packets per second at peak load. The dataset comprised 1.2 million labelled fault events across six fault classes collected over a 90-day monitoring period. All models were trained on an 80/20 stratified split with 5-fold cross-validation.

1. Inference Latency Comparison

Figure 1 presents the mean inference latency (ms) for all evaluated architectures. The proposed Hybrid CNN-LSTM achieves 38.6 ms, remaining below the 50 ms SLA threshold. The simple threshold-based baseline, while lowest in latency (12.4 ms), fails to capture non-linear fault signatures.




Fig. 1: Inference Latency Comparison Across Fault Detection Models (ms). Red dashed line denotes the 50 ms SLA threshold.


2. F1-Score Training Convergence

Fig. 2: F1-Score Convergence Over 50 Training Epochs. Hybrid CNN-LSTM achieves the highest terminal F1 with the least variance.

Figure 2 illustrates F1-score convergence over 50 training epochs for all architectures. The proposed Hybrid CNN-LSTM converges to F1 = 0.963 by epoch 45, exhibiting lower variance and faster convergence than standalone LSTM and BiLSTM models. This is attributed to the convolutional front-end extracting shift-invariant local patterns before temporal modelling.

3. Resource Utilization

Figure 3 compares normalised resource utilisation (CPU, memory, bandwidth) across all models. The Hybrid CNN-LSTM achieves a balanced profile: CPU at 44%, memory at 430 MB, and bandwidth at 4.4 Mbps—representing a 27.9% reduction in memory relative to BiLSTM, while maintaining higher accuracy.






Fig. 3: Normalised Resource Utilization Comparison — CPU (%), Memory, and Bandwidth per model.

4. Cost vs. Performance Trade-off




Fig. 4: Cost vs. Accuracy Trade-off. The arrow indicates the Pareto frontier direction; Hybrid CNN-LSTM dominates all baselines.

Figure 4 plots accuracy against operational cost per node per hour. The Pareto frontier—the curve of models from which no improvement in accuracy can be achieved without increasing cost—passes through the Hybrid CNN-LSTM, confirming it as the dominant solution. No competing model achieves higher accuracy at equivalent or lower cost.

5. Per-Class Fault Detection Rate

Figure 5 shows fault detection rates across all six diesel engine fault classes. The proposed model achieves the highest rates for all classes, particularly for coolant leak detection (96.3%) and injector clogging (95.8%). Bearing degradation detection shows the lowest rate at 91.7%, reflecting the gradual and multi-modal nature of this failure mode.




Fig. 5: Per-Class Fault Detection Rate (%) — Baseline LSTM vs. Proposed Hybrid CNN-LSTM. Percentage improvements annotated above bars.


6. End-to-End Pipeline Throughput and Efficiency

Figure 6 displays throughput (K packets/sec) and pipeline efficiency (%) at each processing stage from sensor ingestion to alert dispatch. The system maintains above 81% end-to-end efficiency at full load, with the largest efficiency drop occurring at the ML inference stage (83.1%)—primarily due to batch assembly overhead in the Azure Container Instance.

Fig. 6: End-to-End Pipeline Throughput (left axis) and Stage Efficiency % (right axis) from Sensor Ingestion to Alert Dispatch.

IV. MODELING APPROACHES FOR AUTONOMOUS FAULT PREDICTION

Data traditionally has been used in a supervised manner for predicting specific failures in diesel engines. For example, in one effort, keyboard classifiers based on the information from dozens of directly connected sensor modalities were designed to predict defects in drivetrain components six minutes ahead. The classifiers were useful in monitoring operational conditions in the diesel engines as undetected failures could lead to major accidents. In another investigation, the vibration signals from diesel engines were



classified into eight conditions including defective conditions of the engine components. For each operating condition, specific prediction models based on kernel extreme learning were built to handle the non-linear mapping from vibration signals to the corresponding probabilities.

Deep learning classifiers have also been implemented for fault classification of heavy diesel engines using multi-source sensor data. A two-stage approach consisting of fault detection and classification was formulated. The first stage predicts whether a fault will occur in the foreseeable future using vibration signals and motor current signals as time series data, while the second stage predicts the type of imminent fault based on wideband sonar signals. The time interval between prediction and occurrence is approximately 2 hours, allowing extensive time for maintenance. In another work, the authors proposed a visual monitoring method for diesel motors by exploiting deep learning techniques based on captured real-life videos. Convnets and recurrent neural networks were implemented in combination to automatically recognize the operational states of diesel engines from video sequences.

More close to the proposed approach, deep learning models also were implemented to closely monitor diesel engines based on limited sensor modalities and on experimental data only. However, their usage was still confined to more classical deep learning-based detection/classification problems that rely on the availability of large training datasets that are usually difficult to obtain.

A. Performance Comparison Tables


1. Comparative Performance of Fault Detection Models

Table 1 provides a comprehensive comparative evaluation of all models across key performance indicators. The proposed Hybrid CNN-LSTM consistently outperforms all baselines on accuracy, F1-score, precision, and recall.

Metric	Threshold Baseline	CNN-1D	LSTM	Hybrid CNN-LSTM (Proposed)	Improv. (%)
Accuracy (%)	76.1	87.3	91.1	96.3	+4.3 vs BiLSTM
F1-Score (macro)	0.734	0.861	0.903	0.963	+5.0%
Precision (macro)	0.748	0.874	0.916	0.971	+4.6%

Recall (macro)	0.721	0.849	0.890	0.955	+5.3%
AUC-ROC	0.801	0.921	0.948	0.984	+2.6%
Miss Rate (%)	27.9	15.1	11.0	4.5	-51.6%
False Alarm Rate (%)	24.2	12.4	8.9	3.2	-58.9%

Table 1: Comparative Performance of Fault Detection Models. Best values in each row are from the Proposed model. Improvement % is relative to the next-best baseline (BiLSTM).



2. Latency and Error Metrics

Table 2 details end-to-end latency decomposition and error characteristics. Latency values are measured as mean ± standard deviation over 10,000 inference cycles on the Azure deployment.


Stage / Metric	Threshold Baseline	CNN-1D	LS TM	BiLSTM	Hybrid CNN-LSTM (Proposed)
Acquisition Latency (ms)	1.2 ± 0.1	1.2 ± 0.1	1.2 ± 0.1	1.2 ± 0.1	1.2 ± 0.1
Transmission Latency (ms)	3.4 ± 0.6	3.4 ± 0.6	3.4 ± 0.6	3.4 ± 0.6	3.4 ± 0.6
Processing Latency (ms)	4.1 ± 0.4	5.8 ± 0.5	6.2 ± 0.6	6.7 ± 0.8	6.1 ± 0.5
Inference Latency (ms)	3.7 ± 0.3	18.3 ± 1.1	36.5 ± 2.4	43.8 ± 3.1	27.9 ± 1.8
Total Latency L_total (ms)	12.4 ± 0.9	28.7 ± 1.6	47.3 ± 2.9	55.1 ± 3.7	38.6 ± 2.2
SLA Compliance (50 ms)	✓	✓	✓	✗	✓
RMSE (fault score)	0.214	0.148	0.119	0.107	0.063
MAE (fault score)	0.178	0.121	0.094	0.083	0.051

Table 2: End-to-End Latency Decomposition and Error Metrics. ✓/✗ indicate SLA compliance at the 50 ms threshold. Values are mean ± std over 10,000 inference cycles.

3. Dataset and Model Architecture Summary

Table 3 summarises the dataset characteristics and architectural hyperparameters of the proposed Hybrid CNN-LSTM model.

Parameter	Value / Configuration
Engine Fleet Size	48 diesel engine nodes
Sensors per Node	12 (vibration, temp., pressure, RPM, EGR, O ₂ ×2, fuel flow ×2, coolant, NO _x)
Sampling	100 Hz per sensor



Fault Classes C	6 (injector, turbo, piston, coolant, bearing, EGR)
Train / Validation / Test Split	72% / 8% / 20% (stratified)
CNN Layers	3 × Conv1D (filters: 64, 128, 256; kernel: 3)
LSTM Layers	2 × Bidirectional LSTM (units: 128, 64)
Fully Connected Layers	2 × Dense (128 → 64) + Softmax output (C=6)
Dropout Rate	0.30 (applied after each LSTM layer)
Optimizer	Adam ($\text{lr} = 1 \times 10^{-3}$, $\beta_1=0.9$, $\beta_2=0.999$)
Regularization λ	1×10^{-4} (L2 weight decay)
Batch Size	256
Training Epochs	50 (early stopping patience = 8)
Total Trainable Parameters	1,847,622
Azure Deployment	Azure Container Instances (ACI), Standard_D2s_v3
IoT Hub Tier	Standard S1 (400,000 msgs/day/unit)
Stream Analytics Streaming Units	6 SUs

Table 3: Dataset Characteristics and Hybrid CNN-LSTM Architecture Hyperparameters.

V. SYSTEM ARCHITECTURE FOR CONNECTED DIESEL ENGINE ECOSYSTEMS

Designing and implementing a decentralized, real-time, and dynamic data monitoring protocol is essential for establishing a system architecture that supports the connected diesel engine ecosystem. Such an architecture provides robust data ingestion and information exchange support by deploying appropriate communication protocols, as the functionality of the fault prediction model is limited by the quality of the input data. Data sender nodes comprising full-duplex Raspberry Pi units are provisioned with a TCP server

to receive the sensor-specific data streams exhibited by auxiliary and environmental sensor subsystem sub-entities. The data sender that harvests these supporting sensor sources executes the data communication in a chained manner considering the data directions. The sub-entities are configured with minimal RAM and computing resource requirements to continue data operations without interruption and only enter sleep mode.

The two actuator nodes comprising set-one and set-two diesel engine actuators inform the IoT read-modules of their triggered events and are controlled by a TCP client in the respective actuator nodes. The main and the TCP receiver-client of the communication subsystem are assigned to the gateway node that operates as the controller and scheduler of the connected diesel engine ecosystem. The controller and the prediction part of the system architecture communicate with

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 02 ISSUE: 04

RECEIVED: OCTOBER 12

REVISED: OCTOBER 28

ACCEPTED: NOVEMBER 09

PUBLISHED: NOVEMBER 20

each other over the user datagram protocol (UDP). The communication subsystem further exploits a transport layer security (TLS) extension over two of the users of this protocol to ensure secure information transmission and reception. The fault prediction system of the diesel engine ecosystem communicates with the image acquisition submodules for visual data sources in an event-driven manner to stop the image-data overflow by activating it only on fault occurrence.

A. Data Ingestion and Communication Protocols

The data ingestion functionality is responsible for receiving and storing data from internal, embedded sensors in the engine control unit (ECU) as well as data from external, Internet-connected devices. This capability makes use of the following protocols:

1. Internally mounted sensors in the ECU that are powered and connected to the ECU's network for data transfer. The controller area network (CAN) is the standard vehicular bus network for connecting these sensors. Data received through the CAN is converted to JSON format and stored in the Temperature_Body_White_Clouds_Made_24/7 collection in a Firestore database.

2. Data received from the external environment that is powered through DC sources. It is connected to a router through the Internet of Things (IoT) communication protocol MQTT (Message Queuing Telemetry Transport) and provided to the cloud. Sensors that are connected and their corresponding input data to the Firestore database collection are as follows: HMC5883L, temperature and humidity sensors, respectively, and DHT, heat index (HI).

3. Data that is fed into the system through external sources, except internal vehicle-mounted sensors, is made available through the 5G network. External sources of data include screenshots of social media posts on different weather parameters such as rainfall, snow, etc. In such cases, Telegram API is used to receive messages. In situations where images are transmitted for prediction, those images are uploaded in an AWS S3 bucket and stored. The URLs of these uploaded files in the S3 bucket are made accessible in the Firestore collection called Image_upload_VEWC.

VI. CONCLUSION

The autonomous fault prediction architecture for connected diesel engine ecosystems discussed in this work relies on a systematic strategy to prepare data for deep learning. An evidence-based selection of model architectures balances accuracy, prediction time, and fault class

distribution. The configuration relies on a private cloud for data ingestion and on-premises meta-cloud servers to address local prediction needs while accelerating the machine-agnostic models. The proposed setup, which considers logistics for test vehicles travelling on picture-perfect road conditions and performs image preprocessing during model training, achieves accurate predictions.

Generated findings and data collection processes serve as a foundation to explore other modalities, including vibration, noise, and air-fuel ratio datasets. Such an approach ensures seamless deployment in the Industrial Internet of Things (IIoT), creating additional value for vehicle fleet operators, Original Equipment Manufacturers (OEMs), and subsequent supply chain producers by enabling e-maintenance solutions.

REFERENCES

1. Theissler, A., Pérez-Velázquez, J., Kettelgerdes, M., & Elger, G. (2021). Predictive maintenance enabled by machine learning: Use cases and challenges in the automotive industry. *Reliability Engineering & System Safety*, 215, 107864.
2. Huang, H., Yang, L., Wang, Y., Xu, X., & Lu, Y. (2021). Digital twin-driven online anomaly detection for an automation system based on edge intelligence. *Journal of Manufacturing Systems*, 59, 138–150.
3. Ravikumar, K. N., Yadav, A., Kumar, H., Gangadharan, K. V., & Narasimhadhan, A. V. (2021). Gearbox fault diagnosis based on multi-scale deep residual learning and stacked LSTM model. *Measurement*, 182, 110099.
4. Srinivas Kalisetty. (2023). Big Data–Driven Cloud Collaboration Models for Enhancing Supplier–Retailer Synchronization in Modern Manufacturing Supply Chains. *Journal of Computational Analysis and Applications (JoCAAA)*, 31(4), 2188–2205. Retrieved from <https://eudoxuspress.com/index.php/pub/article/view/4232>
5. Giordano, D., Pastor, E., Giobergia, F., Cerquitelli, T., Baralis, E., Mellia, M., Neri, A., & Tricarico, D. (2021). Dissecting a data-driven prognostic pipeline: A powertrain use case. *Expert Systems with Applications*, 180, 115109.
6. Markudova, D., Mishra, S., Cagliero, L., Vassio, L., Mellia, M., Baralis, E., Salvatori, L., & Loti, R. (2021). Preventive maintenance for heterogeneous industrial vehicles with incomplete usage data. *Computers in Industry*, 130, 103468.
7. Chen, C., Liu, Y., Sun, X., Di Cairano-Gilfedder, C., & Titmus, S. (2021). An integrated deep learning-based approach for automobile maintenance prediction with GIS data. *Reliability Engineering & System Safety*, 216, 107919.
8. Yu, W., & Liu, Y. (2023). Edge computing-assisted IoT framework with an autoencoder for fault detection in manufacturing predictive maintenance. *IEEE Transactions on Industrial Informatics*, 19(4), 5701–5710.
9. Pamisetty, V. (2023). Transforming Community Engagement with Generative AI: Harnessing Machine Learning and Neural

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 02 ISSUE: 04

RECEIVED: OCTOBER 12

REVISED: OCTOBER 28

ACCEPTED: NOVEMBER 09

PUBLISHED: NOVEMBER 20

- Networks for Hunger Alleviation and Global Food Security. Journal for Re Attach Therapy and Developmental Diversities.
10. Jiang, Y., Dai, P., Fang, P., Zhong, R. Y., Zhao, X., & Cao, X. (2022). A²-LSTM for predictive maintenance of industrial equipment based on machine learning. *Computers & Industrial Engineering*, 172, 108560.
 11. Moin, A., Challenger, M., Badii, A., & Günnemann, S. (2022). Supporting AI engineering on the IoT edge through model-driven TinyML. In *Proceedings of the 2022 IEEE 46th Annual Computers, Software, and Applications Conference* (pp. 884–893). IEEE.
 12. Karimi, E., Chen, Y., & Akbari, B. (2022). Task offloading in vehicular edge computing networks via deep reinforcement learning. *Computer Communications*, 189, 193–204.
 13. Gültekin, Ö., Cinar, E., Özkan, K., & Yazıcı, A. (2022). Real-time fault detection and condition monitoring for industrial autonomous transfer vehicles utilizing edge artificial intelligence. *Sensors*, 22(9), 3208.
 14. Gültekin, Ö., Cinar, E., Özkan, K., & Yazıcı, A. (2022). Multisensory data fusion-based deep learning approach for fault diagnosis of an industrial autonomous transfer vehicle. *Expert Systems with Applications*, 200, 117055.
 15. Reddy, R. (2023). Predictive Health Insights: Ai And Ml's Frontier In Disease Prevention And Patient Management. Available at SSRN.
 16. Gong, C.-S. A., Su, C.-H. S., Liu, Y.-E., Guu, D.-Y., & Chen, Y.-H. (2022). Deep learning with LPC and wavelet algorithms for driving fault diagnosis. *Sensors*, 22(18), 7072.
 17. Gong, C. S. A., Su, C. H. S., Chen, Y. H., & Guu, D. Y. (2022). How to implement automotive fault diagnosis using artificial intelligence scheme. *Micromachines*, 13(9), 1380.
 18. Wu, J.-D., Ke, J.-Y., Shih, F.-Y., & Shyr, W.-J. (2022). Fault diagnosis for vehicle air conditioning blower using deep learning neural network. *Journal of Low Frequency Noise, Vibration and Active Control*, 41(3), 910–925.
 19. Hernández, J. C. M., Madrid, F. G., Quintero, H. F., & Alzate, J. D. R. (2022). Diesel engine diagnosis based on entropy of vibration signals and machine learning techniques. *Electronics Letters*, 58(11), 442–444.
 20. Jiang, J., Li, H., Mao, Z., Liu, F., Zhang, J., Jiang, Z., & Li, H. (2022). A digital twin auxiliary approach based on adaptive sparse attention network for diesel engine fault diagnosis. *Scientific Reports*, 12, 675.
 21. Meda, R., & Pamisetty, A. (2023). Intelligent Infrastructure for Real-Time Inventory and Logistics in Retail Supply Chains. *Educational Administration: Theory and Practice*, 29 (4), 5215–5233.
 22. Yan, H., Bai, H., Zhan, X., Wu, Z., Wen, L., & Jia, X. (2022). Combination of VMD mapping MFCC and LSTM: A new acoustic fault diagnosis method of diesel engine. *Sensors*, 22(21), 8325.
 23. Qin, C., Jin, Y., Zhang, Z., Yu, H., Tao, J., Sun, H., & Liu, C. (2023). Anti-noise diesel engine misfire diagnosis using a multi-scale CNN-LSTM neural network with denoising module. *CAAI Transactions on Intelligence Technology*.
 24. Viana, D. P., de Sá Só Martins, D. H. C., de Lima, A. A., Silva, F., Pinto, M. F., Gutiérrez, R. H. R., Monteiro, U. A., Vaz, L. A., Prego, T., Andrade, F. A. A., Tarrataca, L., & Haddad, D. B. (2023). Diesel engine fault prediction using artificial intelligence regression methods. *Machines*, 11(5), 530.
 25. Cai, Y., Xu, Z., Wen, Q., Shi, J., Zhong, F., Yang, X., Yang, J., & Zhou, H. (2023). Fault states diagnosis of marine diesel engine valve based on a modified VGG16 transfer learning method. *Mathematical Problems in Engineering*, 2023, 1225536.
 26. Pandugula, C., & Nampalli, R. C. R. Optimizing Retail Performance: Cloud-Enabled Big Data Strategies for Enhanced Consumer Insights.
 27. Giannoulidis, A., & Gounaris, A. (2023). A context-aware unsupervised predictive maintenance solution for fleet management. *Journal of Intelligent Information Systems*, 60, 521–547.
 28. Min, H., Fang, Y., Wu, X., Lei, X., Chen, S., Teixeira, R., Zhu, B., Zhao, X., & Xu, Z. (2023). A fault diagnosis framework for autonomous vehicles with sensor self-diagnosis. *Expert Systems with Applications*, 224, 120002.
 29. Asutkar, S., Chalke, C., Shivgan, K., & Tallur, S. (2023). TinyML-enabled edge implementation of transfer learning framework for domain generalization in machine fault diagnosis. *Expert Systems with Applications*, 213, 119016.
 30. Brunheroto, P. H., Pepino, A. L. G., Deschamps, F., & Loures, E. R. (2022). Data analytics in fleet operations: A systematic literature review and workflow proposal. *Procedia CIRP*, 107, 1192–1197.