
Assessment of Road Surface Conditions Using Machine Learning

Krishna Yadav

Assistant Professor

Department of Civil Engineering

Sri Indu College of Engineering & Technology, Telangana

Email Id: *Krishnacivil03@gmail.com*

S.Priyanka¹, P. Guruswamy Goud²

Assistant Professor

Department of Civil Engineering

St. Martin's Engineering College, Telangana

Email Id: *sambari.priyanka93@gmail.com¹, guruswamygoud@gmail.com²*

Abstract

Interstate and asphalt configuration assumes an Assessing road surface conditions is essential for effective transportation infrastructure management, ensuring safety, operational efficiency, and cost-effective maintenance. Conventional evaluation methods are often slow, labor-intensive, and susceptible to human error. This study investigates the use of machine learning techniques to automate and enhance the accuracy of road condition assessments. By integrating image processing, sensor data, and deep learning algorithms, the proposed approach enables real-time detection and classification of pavement distresses such as cracks, potholes, and rutting. The research utilizes data from multiple sources, including satellite imagery, LiDAR (Light Detection and Ranging), and vehicle-mounted sensors, to develop predictive models for road deterioration trends. Findings indicate that machine learning significantly improves the efficiency and reliability of road surface evaluations compared to traditional methods. This study underscores the transformative role of artificial intelligence in modernizing transportation infrastructure maintenance, leading to more sustainable and cost-efficient road management solutions.

Keywords: *Road Surface Assessment, Machine Learning, Pavement Distress Detection, Deep Learning, Artificial Intelligence, Transportation Infrastructure, LiDAR*

INTRODUCTION

General

The assessment of road surface conditions is a critical aspect of transportation infrastructure management. Accurate and timely evaluation of road conditions helps in ensuring traffic safety, reducing vehicle operating costs, and optimizing road maintenance strategies. Traditional assessment methods, such as manual inspections and sensor-based surveys, often involve significant labor, time, and financial resources. These approaches also have limitations in scalability and objectivity, necessitating the need for automated and data-driven solutions.

Recent advancements in machine learning (ML) and computer vision have revolutionized the field of road condition monitoring by enabling efficient, accurate, and real-time assessments. Machine learning models can process vast amounts of data collected from diverse sources, including high-resolution cameras, LiDAR sensors, accelerometers, and GPS-equipped vehicles. These models can detect road surface anomalies such as potholes, cracks, ruts, and uneven surfaces, significantly improving the speed and accuracy of road condition evaluation.

State-of-the-art approaches utilize deep learning techniques, such as Convolution Neural Networks (CNNs) and Transformers, for image-based defect detection, while traditional methods like Support Vector Machines (SVMs) and Random Forest classifiers have been effective in structured datasets. Additionally, unsupervised learning and anomaly detection methods, including auto encoders and clustering algorithms, have been employed for identifying emerging road deterioration patterns.

Research studies have demonstrated the effectiveness of ML in road assessment. For instance, a study by Zhang et al. (2023) showed that CNN-based models trained on aerial and ground-level imagery achieved over 90% accuracy in detecting surface defects. Another study by Gupta et al. (2022) integrated LiDAR and accelerometer data with ML models, achieving high precision in road roughness classification.

This study aims to explore the application of machine learning in road surface condition assessment, covering various data collection techniques, feature extraction methods, and model selection strategies. By leveraging AI-driven approaches, transportation agencies can enhance predictive maintenance, reduce repair costs, and improve overall road safety for both drivers and pedestrians.

METHODOLOGY

The methodology for assessing road surface conditions using machine learning involves multiple steps, including data collection, preprocessing, feature extraction, model selection, training, evaluation, and deployment. The process is designed to ensure accurate, scalable, and real-time assessment of road surface conditions.

Data Collection

The first step is gathering diverse data sources to analyze road surface conditions. These sources include

- **Camera-based Data:** High-resolution images/videos captured by vehicle-mounted cameras or drones.
- Publicly available road condition datasets (e.g., CRACK500, RDD datasets).
- **LiDAR (Light Detection and Ranging) Data:** Captures 3D road surface details using laser-based scanning. Helps in detecting surface irregularities and elevation changes.
- **Accelerometer & Gyroscope Data:** Collected from smart phones or vehicle-mounted sensors to measure road vibrations and roughness. Helps detect bumps, potholes, and cracks.
- **GPS & GIS Data:** Provides geospatial mapping of road defects. Enables location-based road condition assessment.
- **Thermal & Infrared Imaging (Optional):** Used in some cases to detect subsurface road defects.

Data Preprocessing

Since collected data may contain noise or inconsistencies, preprocessing is crucial for improving model performance. The key preprocessing steps include:

- **Image Data Processing**
Resizing, grayscale conversion, and contrast enhancement for better defect visibility.
- **Data augmentation** (flipping, rotation, brightness adjustments) to improve generalization.

Sensor Data Processing

- Noise filtering using Kalman filters or wavelet transforms to smooth accelerometer signals.
- Segmentation of time-series data to analyze specific road segments.

LiDAR Data Processing

- Point cloud processing to remove outliers.
- Surface reconstruction techniques to identify road deformations.

Data Labeling & Annotation

- Manual or semi-automated labeling of defects (e.g., cracks, potholes, ruts).
- Use of pre-labeled datasets if available.

MODEL EVALUATION & VALIDATION

To ensure high accuracy and reliability, the trained models are evaluated using various performance metrics:

For Image-Based Models

- Accuracy, Precision, Recall, F1-score.
- Intersection-over-Union (IoU) for object detection models.

For Sensor-Based Models

- Root Mean Square Error (RMSE), Mean Absolute Error (MAE).
- R-Squared (R^2) for regression-based roughness prediction.

Cross-Validation Techniques

- K-Fold Cross-Validation for improving generalization.
- Leave-One-Out Cross-Validation (LOOCV) for small data sets.

FUTURE ENHANCEMENTS

Future improvements in ML-based road assessment include:

- Federated Learning for privacy-preserving road condition monitoring.
- Multi-Modal Fusion of LiDAR, thermal, and radar data for more robust detection.
- Self-Supervised Learning to reduce dependency on labeled data.
- Integration with Autonomous Vehicles for continuous, real-time road condition monitoring.

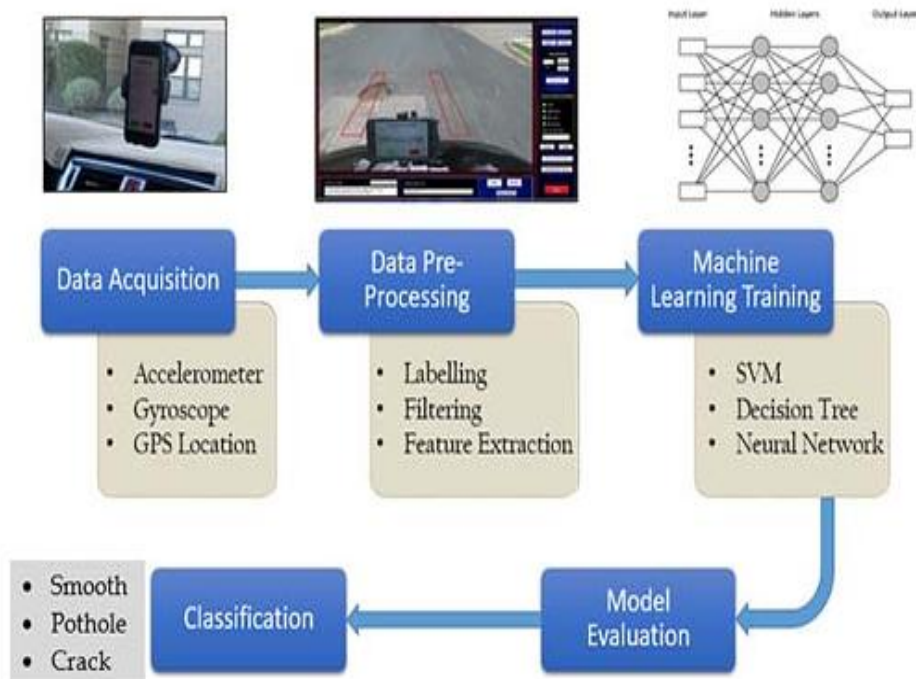


Figure no. 1: Assessing Road Surface Anomaly Using Smart Phone Sensors and Machine Learning

LITERATURE REVIEW

Overview

Before the integration of ML, road assessment relied on manual inspections, sensor-based techniques, and image-processing methods.

- **Manual Inspections:** Traditional methods involved human inspectors visually examining roads for cracks, potholes, and rutting. While effective, these methods were labor-intensive, subjective, and inconsistent (Zhang et al., 2018).
- **Sensor-Based Approaches:** Devices such as accelerometers, gyroscopes, and laser profilers were used to measure road roughness and vibrations (Gonzalez et al., 2019). However, these methods required expensive hardware and complex calibration.
- **Image Processing Techniques:** Early studies leveraged edge detection and texture analysis to detect cracks and potholes in images (Mathavan et al., 2016). While useful, these techniques struggled with lighting variations and occlusions.

RESEARCH METHODS

The research follows a quantitative, data-driven approach, leveraging machine learning models to classify and predict road conditions. The study typically includes

- **Experimental Research:** Models are trained and tested on real-world datasets.
- **Comparative Analysis:** Different ML models are compared to determine the most effective approach.
- **Case Studies & Pilot Studies:** Testing models in specific cities or road networks.

RSC CLASSIFICATION

CNN Model

This category of deep learning models is well known for hybrid feature abstraction from images [22]. While the convolution (C) layers extract neighborhood pixel contributions, the sub-sampling (S) /pooling layers followed by a non-linear activation function enable multi-resolution feature learning. A sample architecture of a CNN model with 3 convolution/sub-sampling layers (CS3) followed by a dense (D) flattening layer. Optimal parameterization for the CNN model involves selection of sufficient C-S layers, and optimal filter bank sizes to ensure low trainable parameters, thereby preventing over fitting. Often, deep CNN models with many hidden layers of batch normalization [19] makes the network train faster while ensuring low training losses.

The numbers of filters in each convolution layer double in each consecutive stage. Thus, for n_T training images, and n_c number of output classes, the cross-entropy loss function used for training NN-based models by back-propagation is given in (1),

$$L = - \sum_{i=1}^{n_T} \sum_{l=1}^{n_c} \mathbb{1}_{[t_i = l]} \log(p_{sl}^i) \quad (1)$$

where, $\mathbb{1}_{[t_i = l]}$ is the indicator function that GT vector t_i is equal to class label l , while p_{sl}^i represents the model predicted probability that image i belongs to class label l . The ReLU activation function applied on weighted sum of outputs from the previous layer i_0 are given in (2).

$$y_{i'} = \sum_{k=1}^n x_k w_{ki'}, \quad p_{si'}^i = \max(0, y_{i'}) \quad (2)$$

Here, n signifies the number of neurons present in the layer just before the activation layer. The final predicted class label is the class with maximum predicted probability in (3).

$$s_i = \arg_l \max(p_{sl}^i) \quad \forall l \text{ classes.} \quad (3)$$

Table no 1: Feature Extraction Average Time Requirements

Parameter	Using Features from all Axes (ms)	Using Features from Y' Axis (ms)
High Pass Filtering	0.0157	0.0157
Time Domain Feature Extraction	15.257	5.135
Frequency Domain Feature Extraction	1.674	1.084
Wavelet Domain Feature Extraction	52.877	19.01
Total	69.823	25.245

DATA PREPROCESSING METHODS

Since raw data is often noisy and inconsistent, preprocessing is necessary to improve model performance.




Figure no 2: Cartesian coordinate Axes of Phone Accelerometer and Gyroscope

Image Data Preprocessing

- **Noise Reduction:** Gaussian filtering, histogram equalization.
- **Data Augmentation:** Image flipping, rotation, brightness adjustment to increase dataset variability.
- **Segmentation & Annotation:** Manually or automatically labeling cracks, potholes, and surface anomalies.

Sensor Data Preprocessing

- **Noise Filtering:** Applying Kalman filters or wavelet transforms to remove unwanted vibrations.
- **Time-Series Segmentation:** Dividing data into meaningful segments for classification.

LiDAR Data Processing

- **Point Cloud Processing:** Removing outliers and filtering noise.
- **3D Surface Reconstruction:** Creating elevation maps for defect detection.

FEATURE EXTRACTION & SELECTION

Feature extraction helps machine learning models differentiate between different road conditions.



Figure no.3: Cartesian Coordinate Axes of Phone Accelerometer and Gyroscope

Image-Based Feature Extraction

- **Texture Analysis:** Using Gabor filters, Histogram of Oriented Gradients (HOG).
- **Deep Feature Extraction:** Using pre-trained CNNs (ResNet, VGG) to extract hierarchical road surface patterns.

Sensor-Based Feature Extraction

- **Statistical Features:** Mean, variance, skewness of accelerometer data.
- **Frequency Domain Features:** Fast Fourier Transform (FFT) for vibration analysis.

Feature Selection Methods

- **Principal Component Analysis (PCA):** Reducing data dimensionality while preserving important information.
- **Recursive Feature Elimination (RFE):** Selecting the most relevant features for classification.

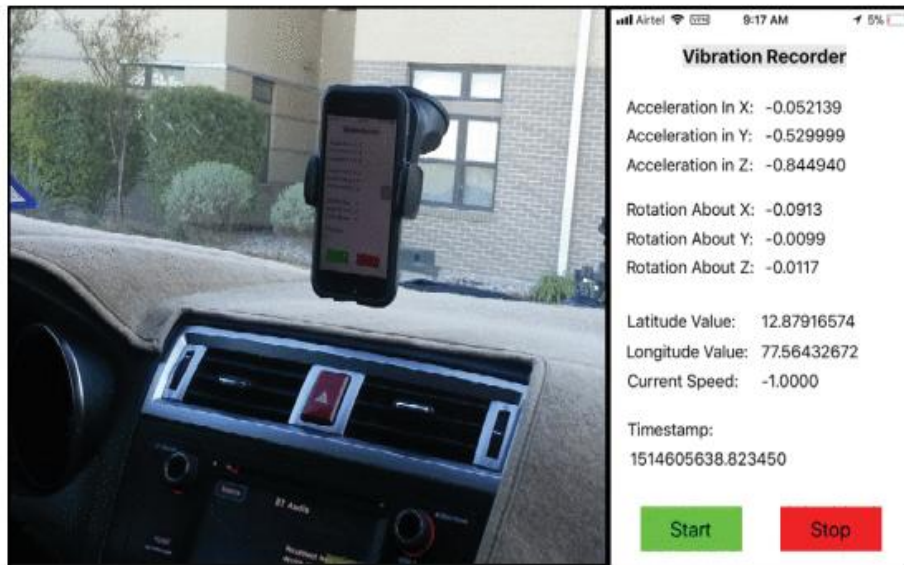


Figure no. 4: Setup of the Phone and Screenshot of the Vibration Recorder App

SELECTING AND TRAINING ML MODELS TO CLASSIFY ROAD SURFACE CONDITIONS

Model Selection

- **For Image-Based Detection**
CNNs (ResNet, YOLO, U-Net) for segmentation & classification.
Vision Transformers (ViTs) for defect recognition.
- **For Sensor-Based Classification**
Support Vector Machines (SVMs), Decision Trees, Random Forests.
Long Short-Term Memory (LSTM) for time-series analysis.
- **For Multi-Modal Learning**
Combining image, sensor, and LiDAR data for enhanced performance.

MODEL TRAINING

- **Dataset Splitting:** Train (70%), Validation (15%), Test (15%).
- **Hyper parameter Tuning:** Grid Search, Bayesian Optimization.
- **Transfer Learning:** Using pre-trained models to improve accuracy.

The reorientation algorithm performs accelerometer data reorientation using Euler's angles, which form a representation of the spatial orientation of a certain reference frame as a combination of three orthogonal elemental rotations. Ideally, when a car is at rest on a flat surface, the acceleration values would be.

$$a_x = 0\text{m/s}^2, \quad a_y = 9.81\text{m/s}^2 \text{ and } a_z = 0\text{m/s}^2$$

Equations (1) to (4) are used to calculate two of the three Euler angles and reorient acceleration values to the global frame. a'_x , a'_y , a'_z are the acceleration values with respect to the global reference frame while α and β are the roll and pitch angles, respectively. Fig. 6. shows the plot of the acceleration data of a 1.5s window before and after reorientation.

$$\alpha = \tan^{-1}\left(\frac{a_y}{a_z}\right), \quad \beta = \tan^{-1}\left(\frac{-a_x}{\sqrt{(a_y)^2 + (a_z)^2}}\right) \quad (1)$$

$$a'_x = \cos(\beta)a_x + \sin(\beta)\sin(\alpha)a_y + \cos(\alpha)\sin(\beta)a_z \quad (2)$$

$$a'_y = \cos(\alpha)a_y - \sin(\alpha)a_z \quad (3)$$

$$a'_z = -\sin(\beta)a_x + \cos(\beta)\sin(\alpha)a_y + \cos(\alpha)\cos(\beta)a_z \quad (4)$$

DESIGN PROCESS

Identify the need for automated road condition monitoring to replace traditional manual inspections. Determine the key road surface defects to detect, such as:

- Cracks (longitudinal, transverse, alligator cracks)
- Potholes
- Rutting
- Roughness (IRI – International Roughness Index)
- Surface Wear

System Objectives

- Improve accuracy in road condition detection.
- Enable real-time analysis for smart cities and road maintenance teams.
- Develop an automated, scalable, and cost-effective monitoring solution.

System Components

An ML-based road assessment system typically includes:

- **Data Acquisition Module:** Captures images, videos, and sensor data (LiDAR, accelerometer, GPS).
- **Data Processing Unit:** Filters and preprocesses raw data for model training.
- **ML Model Training & Classification Module:** Uses AI models to classify and detect road defects.
- **Visualization & Reporting Dashboard:** Displays results in GIS systems or mobile applications for real-world use.

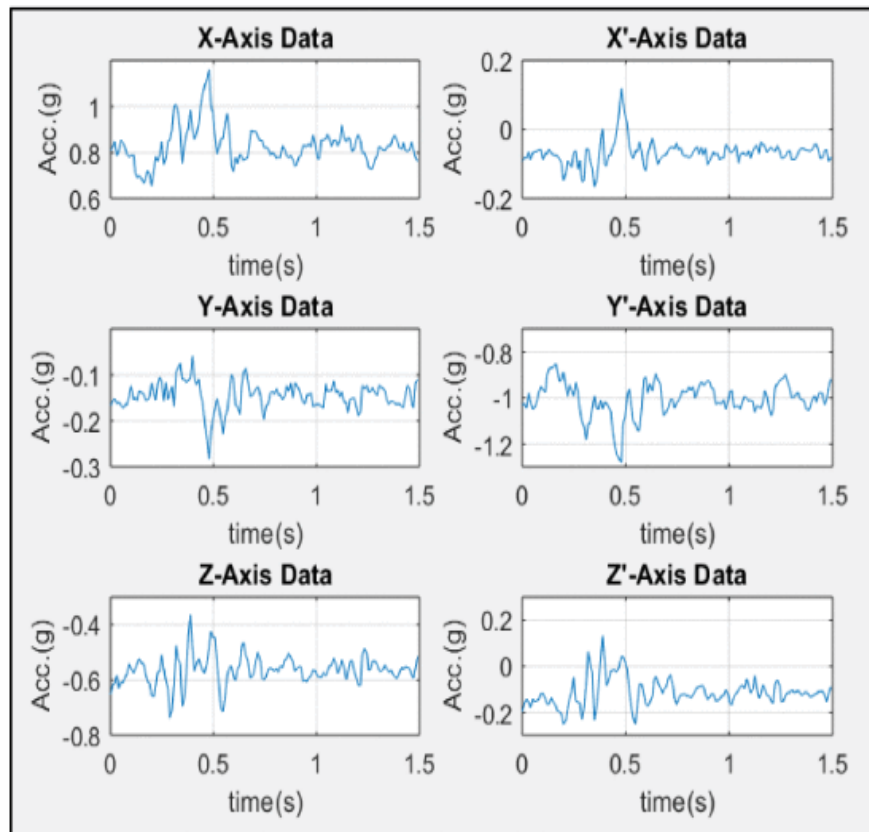


Figure no. 5: Reorientation of acceleration data to global frame of reference

System Workflow

- Data Collection (Images & Sensor Data)
- Data Preprocessing (Noise Removal & Augmentation)
- Feature Extraction & Selection

- ML Model Training & Validation
- Model Deployment (Cloud/Edge Computing)
- Visualization & Decision-Making




Figure no. 6: Road Condition Classifier Software with Tire-Trajectory Overlay

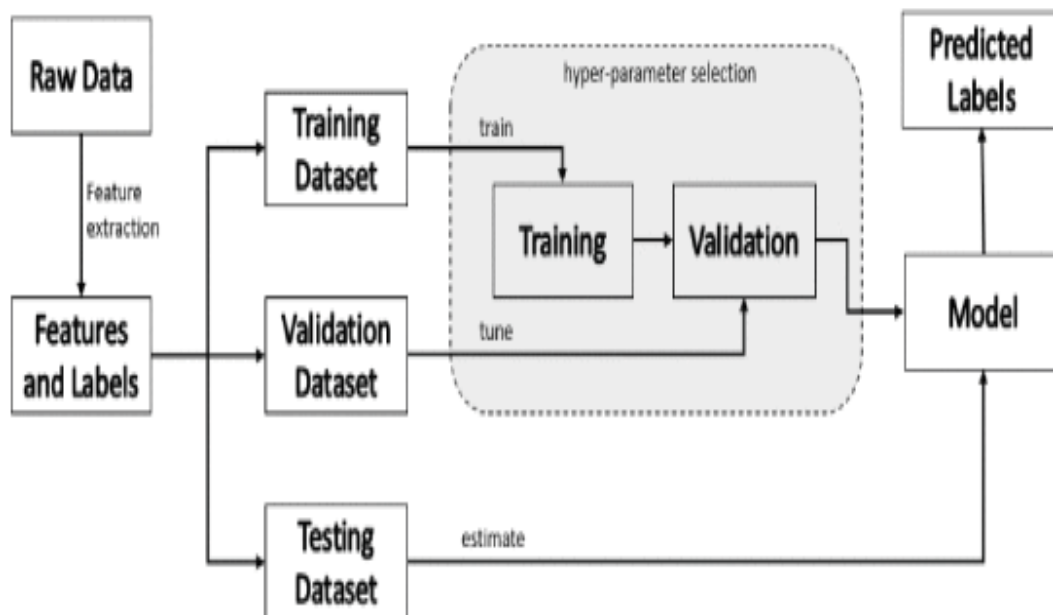


Figure no.7: General Workflow Diagram of Machine Learning Algorithms

FEATURE-BASED MODEL

In this model, the pre-processing steps involve a horizon detector to separate the sky and drivable surface components. For this purpose, GIST feature based horizon detector invoked. Here, the relative horizontal location of the horizon is evaluated as a regression problem using GIST features from the training set of images on the Label Me data set. Next, the sub image under the horizon (R b) is segmented into the drivable surface (R bd) and surroundings (R bs). The sub-image above the horizon (Ra) is similarly segmented into sky (R ak) and surroundings (R as), as shown in Fig. 3. Further, the RGB image plane is transformed to HSV plane, and HOG descriptors are computed for each of the 4 segmented ROIs (R bd, R bs, R ak, R as), to extract color and texture-based features, as shown in Fig. 6. Next, the following 15 color and texture based regional features are extracted per segmented ROI: mean and standard deviation of pixel intensities segmented ROI RGB, HSV plane, and HOG feature planes, and the fraction of pixels belonging to each segmented ROI. The HOG features are extracted using an empirically selected block size of [16x16]. These 15 features corresponding to 4 segmented ROIs result in a 60 feature-vector corresponding to X_i for each image i , that are then subjected to classification using a single hidden layer NN-model. The number of hidden neurons is varied between [10:500], to identify the optimal number of hidden neurons that maximize test accuracy.

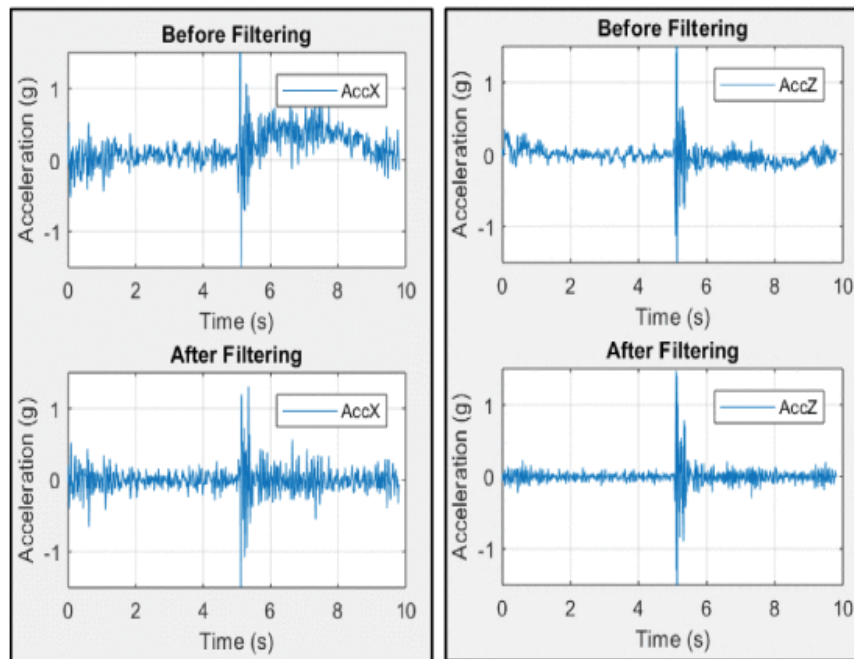


Figure no 8: Acceleration Signal in X' And Z' Axis Before and After Filtering

Examples of road patch quantization and RFE classification. (Top :) Probability score ' $p^i s^1$ ' for patch composition. (Bottom :) Quantized Patches and their GT. (Left :) High RFE for

well-ploughed road. (Center :) Medium RFE for partially patchy road. (Right :) Low RFE for patchy snow and dry surfaces that can lead to slippery roads.

DECISION TREES

Decision trees, also known as classification trees or regression trees, predict output responses based on input data. Following the decisions in the tree from the root node to the leaf node gives the output response to that particular input data. The decision tree is an algorithm that classifies data through a cascade of statistical tests as shown in Fig. 11. These tests compare the value that is input to a node with a threshold value that splits the tree's path. Tests can have multiple results and different tree paths can follow to the same output class label. The complexity of the tree is defined by the number of branch splits and depending on its complexity. Decision trees have quick training and prediction speeds, a moderate predictive accuracy and low computational memory requirements.




Figure no.9: Decision Tree Structure Used in the Work

Neural Networks

Neural networks represent a popular machine learning framework that attempts to imitate the learning pattern of natural biological neural networks in the brain. A typical neural network consists of inter-connected arithmetic processors called neurons which produce a sequence of real valued activation outputs. Neurons present in the input layer of the neural network are activated through sensor data perceiving the environment, while neurons present in other layers get activated through weighted connections from previously active neurons. Neural network algorithms link the feature vectors (input layer) to the class labels (output layer) using multilayered networks called hidden layers, as shown in Fig. 12. The complexity of the

classification problem determines the number of hidden layers needed. Although neural networks are powerful, high accuracy algorithms, training them requires a large dataset. The size of the required dataset also increases as the number of hidden layers increases.




Figure no. 10: Structure of A Neural Network Classifier With Two Hidden Layers. In This Work, We Use 7 To 8 Hidden Layers

DISCUSSIONS

The application of machine learning (ML) in assessing road surface conditions has proven to be an efficient and scalable solution for improving transportation safety, optimizing maintenance efforts, and supporting smart city initiatives. By leveraging sensor data, satellite imagery, and crowd sourced inputs, ML models can provide real-time and predictive insights into road deterioration.

Despite challenges such as data quality, model generalization, and computational constraints, advancements in deep learning, federated learning, and edge computing are enhancing the accuracy and applicability of these systems. The integration of ML with IoT and autonomous vehicle technologies further strengthens its role in modern infrastructure management.

Future research should focus on developing more robust, adaptive, and interpretable ML models to improve reliability across diverse environments. With continued innovation, ML-

driven road condition assessment can significantly enhance road safety, reduce maintenance costs, and support sustainable urban mobility.

FUTURE SCOPE

The future of machine learning (ML) in road surface condition assessment is promising, with several advancements expected to enhance accuracy, scalability, and real-time decision-making. The future of road surface condition assessment using ML will be driven by innovations in data collection, advanced ML models, and integration with smart city infrastructure. With continued research and technological advancements, ML will play a pivotal role in ensuring safer, more efficient and sustainable road networks worldwide.

- Advancements in Data Collection and Integration.
- Improved Machine Learning Techniques.
- Enhanced Predictive and Preventive Maintenance.
- Smart City and Autonomous Vehicle Integration.
- Climate and Disaster Response Applications.

ADVANCEMENTS IN DATA COLLECTION AND INTEGRATION

- **Edge Computing and IoT Integration:** Smart sensors embedded in roads, vehicles, and infrastructure will provide real-time data, reducing reliance on centralized processing.
- **High-Resolution Remote Sensing:** Improved satellite and drone imagery will enhance large-scale road monitoring.
- **Crowd sourced Data from Connected Vehicles:** Autonomous and connected vehicles will continuously feed ML models with up-to-date road condition data.

IMPROVED MACHINE LEARNING TECHNIQUES

- **Self-Supervised and Few-Shot Learning:** Reducing dependence on large labeled datasets will make ML models more adaptable and scalable.
- **Federated Learning:** Decentralized ML models will enable real-time road condition assessments while preserving data privacy.
- **Hybrid AI Models:** Combining deep learning with traditional statistical models and physics-based simulations will improve predictive accuracy.

ENHANCED PREDICTIVE AND PREVENTIVE MAINTENANCE

- **AI-Powered Predictive Maintenance:** ML models will predict road deterioration trends, allowing proactive repair scheduling.
- **Autonomous Road Repair Systems:** Robots and automated maintenance solutions will use ML insights to fix minor road defects before they worsen.

CONCLUSION

The assessment of road surface conditions using machine learning has proven to be a powerful and efficient approach for enhancing road maintenance and safety. By leveraging various data sources such as sensor readings, satellite imagery, and vehicular data, machine learning models can accurately classify and predict road conditions, enabling proactive maintenance strategies.

The study highlights that machine learning techniques, including deep learning, support vector machines, and random forests, offer significant advantages over traditional inspection methods in terms of speed, cost-effectiveness, and accuracy. However, challenges such as data quality, model interpretability, and real-time implementation must be addressed to maximize their effectiveness.

Future research should focus on improving model generalization across diverse geographical regions, integrating real-time data streams, and enhancing explainability to increase adoption by transportation authorities. With continuous advancements in artificial intelligence and data collection technologies, machine learning-based road condition assessment will play a crucial role in building smarter, safer, and more sustainable transportation systems.

REFERENCES

1. U. D. of Transportation Federal Highway Administration, "How do weather events impact roads?" Feb 2017. [Online]. Available: <https://ops.fhwa.dot.gov/weather/q1/roadimpact.htm>
2. K. Koro sec, "2016 was the deadliest year on American roads in nearly a decade," Feb 2017. [Online]. Available: <http://fortune.com/2017/02/15/traffic-deadliest-year/> [3] M. Anderson, F. Bruzelius, J. Casselgren, M. Hjort, S. Lofving, " G. Olsson, J. Ronnberg,

- M. Sjödahl, S. Solyom, J. Svendenius et al., “Road friction estimation part ii,” 2010.
3. S. Khaleghian, A. Emami, and S. Taheri, “A technical survey on tire-road friction estimation,” *Friction*, pp. 1–24, 2017.
 4. I.-H. Kim, J.-H. Bong, J. Park, and S. Park, “Prediction of drivers intention of lane change by augmenting sensor information using machine learning techniques,” *Sensors*, vol. 17, no. 6, p. 1350, 2017.
 5. A. Raj, D. Krishna, K. Shantanu et al., “Vision based road surface detection for automotive systems,” in *Applied Electronics (AE), 2012 International Conference on*. IEEE, 2012, pp. 223–228.
 6. J. H. Kim and J. M. Won, “A development of the road surface decision algorithm using svm (support vector machine) clustering methods,” *The Journal of The Korea Institute of Intelligent Transport Systems*, vol. 12, no. 5, pp. 1–12, 2013.
 7. L. Hansung, S. Choi, and J. Yang, “Driver assistance apparatus capable of recognizing a road surface state and vehicle including the same,” Mar. 21 2017, US Patent 9,598,088.
 8. Q. Zhao, Q. Zhang, J. Wang, B. Litkouhi, and W. Zhang, “Systems and methods for determining a condition of road surface,” Mar. 21 2017, US Patent 9,598,087 B2.
 9. Q. Zhao, J. Wang, W. Zhang, and B. B. Litkouhi, “Road surface condition detection with recursive adaptive learning and validation,” Sep. 22 2015, US Patent 9,139,204.
 10. J. Zhao, H. Wu, and L. Chen, “Road surface state recognition based on svm optimization and image segmentation processing,” *Journal of Advanced Transportation*, vol. 2017, 2017.
 11. S. Roychowdhury and J. Ren, “Non-deep cnn for multi-modal image classification and feature learning: An azure-based model,” in *IEEE International Conference on Big Data (Big Data)*. IEEE, 2016, pp. 2893–2812.
 12. Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
 13. F. N. Iandola, S. Han, M. W. Moskewicz, K. Ashraf, W. J. Dally, and K. Keutzer, “Squeezenet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size,” arXiv preprint arXiv:1602.07360, 2016.

14. A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
15. G. Liu, P. Han, Y. Niu, W. Yuan, Z. Lu, and J.-R. Wen, "Graph boosted convolutional neural networks for semantic segmentation," in *International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2017, pp. 612–618.
16. A. Torralba and P. Sinha, "Statistical context priming for object detection," in *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, vol. 1. IEEE, 2001, pp. 763–770.
17. J. Sivic, B. Kaneva, A. Torralba, S. Avidan, and W. T. Freeman, "Creating and exploring a large photorealistic virtual space," in *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW'08. IEEE Computer Society Conference on*. IEEE, 2008, pp. 1–8.
18. B. C. Russell, A. Torralba, K. P. Murphy, and W. T. Freeman, "Labelme: a database and web-based tool for image annotation," *International journal of computer vision*, vol. 77, no. 1, pp. 157–173, 2008.
19. M. Menenberg, S. Pathak, H. P. Udyapuram, S. Gavirneni, and S. Roychowdhury, "Topic modeling for management sciences: A network based approach," in *2016 IEEE International Conference on Big Data (Big Data)*. IEEE, 2016, pp. 3509–3518.